nature reviews cancer

https://doi.org/10.1038/s41568-026-00922-2

Review article

Check for updates

Advancing AI for multi-omics and clinical data integration in basic and translational cancer research

Fei Liu[1,2], Stephan Beck[3], Lei Yang[4], Huiyan Luo[5] ✉ & Kang Zhang[1,6,7] ✉

Abstract

The extensive heterogeneity of cancer across biological scales necessitates a holistic approach beyond single-analyte methods. Integrating multi-omics data — from genomics to proteomics — with multimodal information, such as clinical records and medical imaging, offers a comprehensive, systems-level view of tumorigenesis. Artificial intelligence (AI) has emerged as the essential technology to decipher these complex, high-dimensional datasets, powering substantial advances in early diagnosis, precise patient stratification, prediction of therapeutic response and the elucidation of mechanisms of drug resistance. To translate these powerful predictive models into practice, explainable AI is critical for building clinical trust and generating novel, testable biological hypotheses. While challenges in data accessibility and model generalizability persist, the field is advancing toward patient-specific digital twins, promising to simulate individual disease trajectories and optimize treatments, thereby heralding a new era of precision oncology.

Sections

[1]Artificial Intelligence Cross Disciplinary Research Institute and Faculty of medicine, Macau University of Science and Technology, Macau, China. [2]National Cancer Center/National Clinical Research Center for Cancer/Cancer Hospital, Chinese Academy of Medical Sciences and Peking Union Medical College, Beijing, China. [3]University College London, Cancer Institute, London, UK. [4]Department of Orthopedics and Key Laboratory of Hepatosplenic Surgery of Ministry of Education, the First Affiliated Hospital, State Key Laboratory of Frigid Zone Cardiovascular Diseases (SKLFZCD), Harbin Medical University, Harbin, China. [5]Department of Medical Oncology, Sun Yat-sen University Cancer Center, State Key Laboratory of Oncology in South China, Guangdong Provincial Clinical Research Center for Cancer, Sun Yat-sen University, Guangzhou, China. [6]Clinical Data Science Institute, Institute for Advanced Study on Eye Health and Diseases, Wenzhou Medical University, Wenzhou, China. [7]Guangzhou National Laboratory, Guangzhou, China. ✉e-mail: luohy@sysucc.org.cn; kang.zhang@gmail.com

Introduction

The advent of high-throughput 'omics' technologies — spanning genomics, epigenomics, transcriptomics, proteomics and metabolomics — has revolutionized our ability to capture comprehensive snapshots of biological systems at an unprecedented depth[1–3]. These technologies generate vast amounts of data, offering the potential to unravel the molecular underpinnings of cancer[4–6]. However, relying on a single-analyte approach can provide an incomplete and potentially misleading picture of tumour biology and patient prognosis[1,7]. For instance, genomic alterations do not always translate directly into functional changes at the protein level, and cellular morphology captured by pathology may not fully reflect the underlying molecular signalling state[4,8].

Multi-omics integration seeks to overcome these limitations by integrating the diverse data layers, alongside clinical information and medical imaging (including radiomics and digital pathology), to provide a more complete picture of the tumour ecosystem and its interactions[1,7,9–11] (Fig. 1). This synergistic approach allows for a more comprehensive characterization of tumour heterogeneity, the interplay between different molecular layers (for example, how genomic changes affect transcription and protein levels), and the connection between molecular profiles, tissue morphology and clinical phenotypes[12–14]. Integrating molecular data with imaging or pathology can bridge the gap between genotype and phenotype by revealing how molecular alterations manifest visually[9,13,15–17]. Similarly, combining omics data with clinical records enables the discovery of molecular correlates of treatment response or resistance observed in patients[1,18,19]. Crucially, in this Review, we adopt a broad, systems-level definition of 'multi-omics' that extends beyond traditional molecular layers to encompass 'digital phenomics' derived from medical imaging (radiomics and pathomics) and codified clinical data. While we acknowledge the emerging significance of the microbiome and exposome in cancer biology, our discussion focuses primarily on the integration of these host-intrinsic molecular, morphological and clinical data streams, as they currently represent the primary pillars for constructing patient-specific digital twins. The sheer scale and high dimensionality of multi-omics datasets

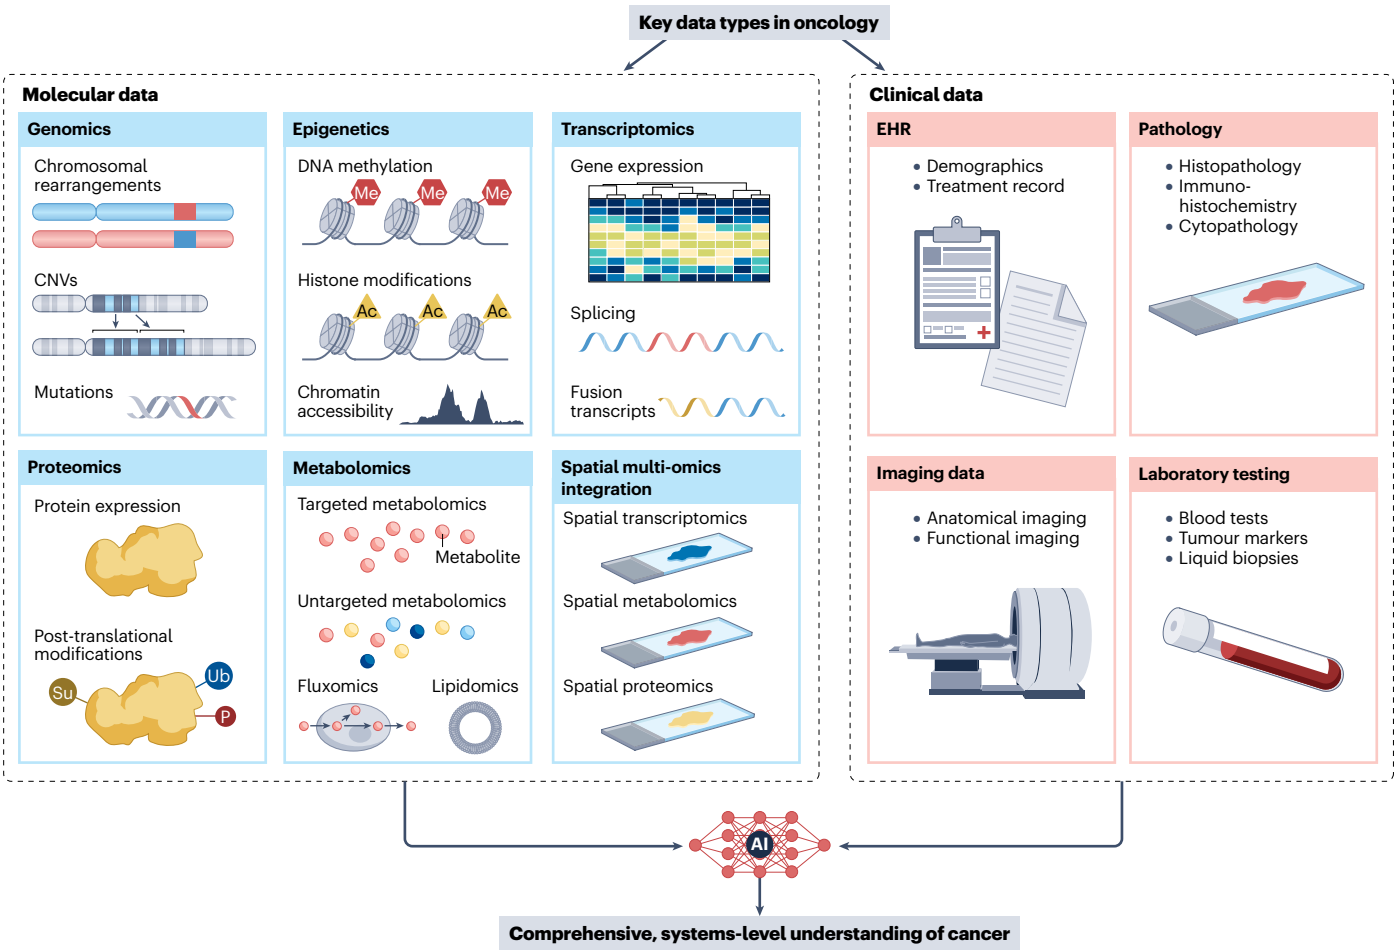

Fig. 1 | The multi-omics data foundation for precision oncology. Modern cancer research is built upon a diverse landscape of multi-omics data that provides a deep molecular characterization of tumours. This includes multiple layers of biological information, from genomic alterations (mutations and copy number variations (CNVs)) and epigenetics (DNA methylation) to transcriptomics, proteomics and metabolomics. These data can be captured at the bulk or single-cell level, and can also be spatially resolved to link molecular profiles to tissue architecture. To translate these rich molecular insights into clinical practice, multi-omics data can be integrated with multimodal clinical data, which provides essential patient context. This includes information from electronic health records (EHRs), medical imaging, pathology and laboratory testing. Artificial intelligence (AI) can then be used to integrate these vast and varied data streams to build a comprehensive, systems-level understanding of an individual patient's cancer, thereby enabling true precision medicine.

present sizable analytical hurdles[4,20]. Integrating disparate data types, each with unique characteristics, noise profiles and biological relevance, requires sophisticated computational methods capable of handling complexity and extracting biologically meaningful insights[1,7]. Compounding these inherent challenges are additional practical hurdles related to data harmonization, particularly when integrating data from various research centres. Variations in technical protocols, such as imaging scanner parameters or sequencing platforms, can introduce considerable batch effects. Furthermore, inter-observer variability in pathological and radiological interpretation adds another layer of complexity that must be addressed. Traditional statistical methods often struggle to model the intricate, nonlinear relationships inherent in such high-dimensional biological data.

Artificial intelligence (AI), particularly deep learning, has emerged as a powerful enabling technology to address these analytical challenges[7,21–26]. AI algorithms excel at identifying complex patterns, integrating diverse data sources and building predictive models from large datasets[7,22,23,27,28]. In the context of multi-omics, AI can facilitate data fusion, uncover hidden correlations between different molecular layers and clinical outcomes, and generate hypotheses about underlying biological mechanisms[1,7,21,29,30]. Crucially, AI models can also be designed to handle the pervasive issue of missing data; it is not always necessary for every patient to have a complete set of all omics modalities. Advanced techniques can learn from incomplete datasets, impute missing values or weigh the available data appropriately, which is vital for leveraging real-world clinical cohorts where data are often heterogeneously collected[31,32]. This flexibility enables the application of a wide range of AI techniques in oncology. For instance, supervised learning is now used to build models that can classify tumour subtypes or predict patient outcomes[33–36], whilst unsupervised methods are essential for discovering novel patient subgroups from integrated data[20]. These powerful computational approaches form the computational backbone for translating multi-omics data into tangible advances in precision oncology.

This Review synthesizes the current state of AI-driven multi-omics data integration in oncology to articulate a new paradigm. While seminal perspectives have established the conceptual frameworks for multimodal data fusion[37], here, we synthesize the rapid evolution of the field driven by next-generation technologies — specifically generative AI, foundation models and digital twins. It explores how this AI-driven synthesis is not merely improving isolated tasks but is creating a cohesive, systems-level understanding of the individual patient. This holistic view is poised to redefine the entire cancer care continuum, from enhancing the elucidation of foundational biological mechanisms to enabling the future of dynamic disease monitoring. To this end, the discussion encompasses the core computational methodologies and the critical role of explainable AI (XAI) in bridging the gap between predictive models and biological plausibility. Ultimately, by weaving together these disparate data threads, multi-omics integration holds the promise of constructing more accurate diagnostic and prognostic models, identifying robust biomarkers, and uncovering novel therapeutic targets that would remain hidden within the confines of single-omics analyses.

AI methodologies for multi-omics integration

Translating the vast landscape of multi-omics data into clinical utility hinges on a sophisticated computational framework. The power of AI is not derived from a single algorithm but from a strategic orchestration of learning paradigms, specialized model architectures and data fusion strategies (Fig. 2). Understanding these core methodological components is essential to appreciating both the potential and complexity of applying AI in oncology[38].

Unsupervised learning

One of the greatest practical challenges in medical AI is the scarcity of large, expertly annotated patient datasets required for training deep learning models from scratch. To overcome this bottleneck, the field is increasingly adopting advanced training paradigms that leverage abundant unlabelled data (Fig. 2a). The dominant approach is unsupervised or self-supervised learning[20,39,40]. In this paradigm, a large model is first trained on a massive corpus of unlabelled data (for example, millions of pathology images or publicly available genomic sequences). The benefit of this initial step is that the model learns fundamental, general-purpose features of the data (for example, the typical morphology of cells in pathology images or common sequence motifs in genomics) without needing any expert annotation. The model learns by solving an 'auxiliary' or 'pretext' task that does not require manual labels such as reconstructing a corrupted input or predicting the relationship between different data fragments[41,42]. This process forces the model to learn rich, generalizable representations of the data's intrinsic structure[43]. This powerful pre-trained model then serves as a highly informed starting point.

For the 'primary task' of clinical interest, the model is then fine-tuned using a much smaller, more specific and labelled dataset. Here, the 'labels' are the ground-truth clinical outcomes of interest, such as patient survival or treatment response, which are expensive and time-consuming to acquire. To be effective, the unlabelled data for pre-training should be of the same modality as the data used for the final task (for instance, pre-training on a large cohort of unlabelled immunohistochemistry (IHC) images to prepare for a task on a smaller, labelled IHC cohort). In the context of multi-omics integration, this paradigm is particularly powerful. Modality-specific encoders can be independently pre-trained on large, unimodal, unlabelled datasets (for example, an image encoder on millions of pathology slides and a sequence encoder on a genomic database). These 'expert' encoders are then combined and jointly fine-tuned on a smaller, but complete, multimodal dataset with ground-truth labels to learn the critical cross-modal relationships.

This two-stage approach has given rise to foundation models, which represent a fundamental shift from task-specific architectures to unified, generalizable platforms. Unlike traditional models trained for a single narrow purpose, foundation models pre-trained on massive datasets, such as Virchow[44] for computational pathology or Evo[43,45] for genomic sequences, learn universal representations of biological data. This enables powerful capabilities such as zero-shot inference, where a model can perform a task it was not explicitly trained for, and data-efficient fine-tuning, allowing high performance on rare cancer subtypes with very few labelled examples. Furthermore, emerging multimodal foundation models facilitate cross-modal alignment, mapping distinct data types into a shared latent space that captures their underlying biological meaning[46]. This allows for complex reasoning tasks, such as generating diagnostic reports directly from images, thereby moving AI from a pattern-recognition tool to a comprehensive clinical assistant[47].

Deep learning architectures

Integrating the diverse data types inherent to multi-omics requires a corresponding diversity of deep learning architectures as each is specialized to process a different form of information. The core of any

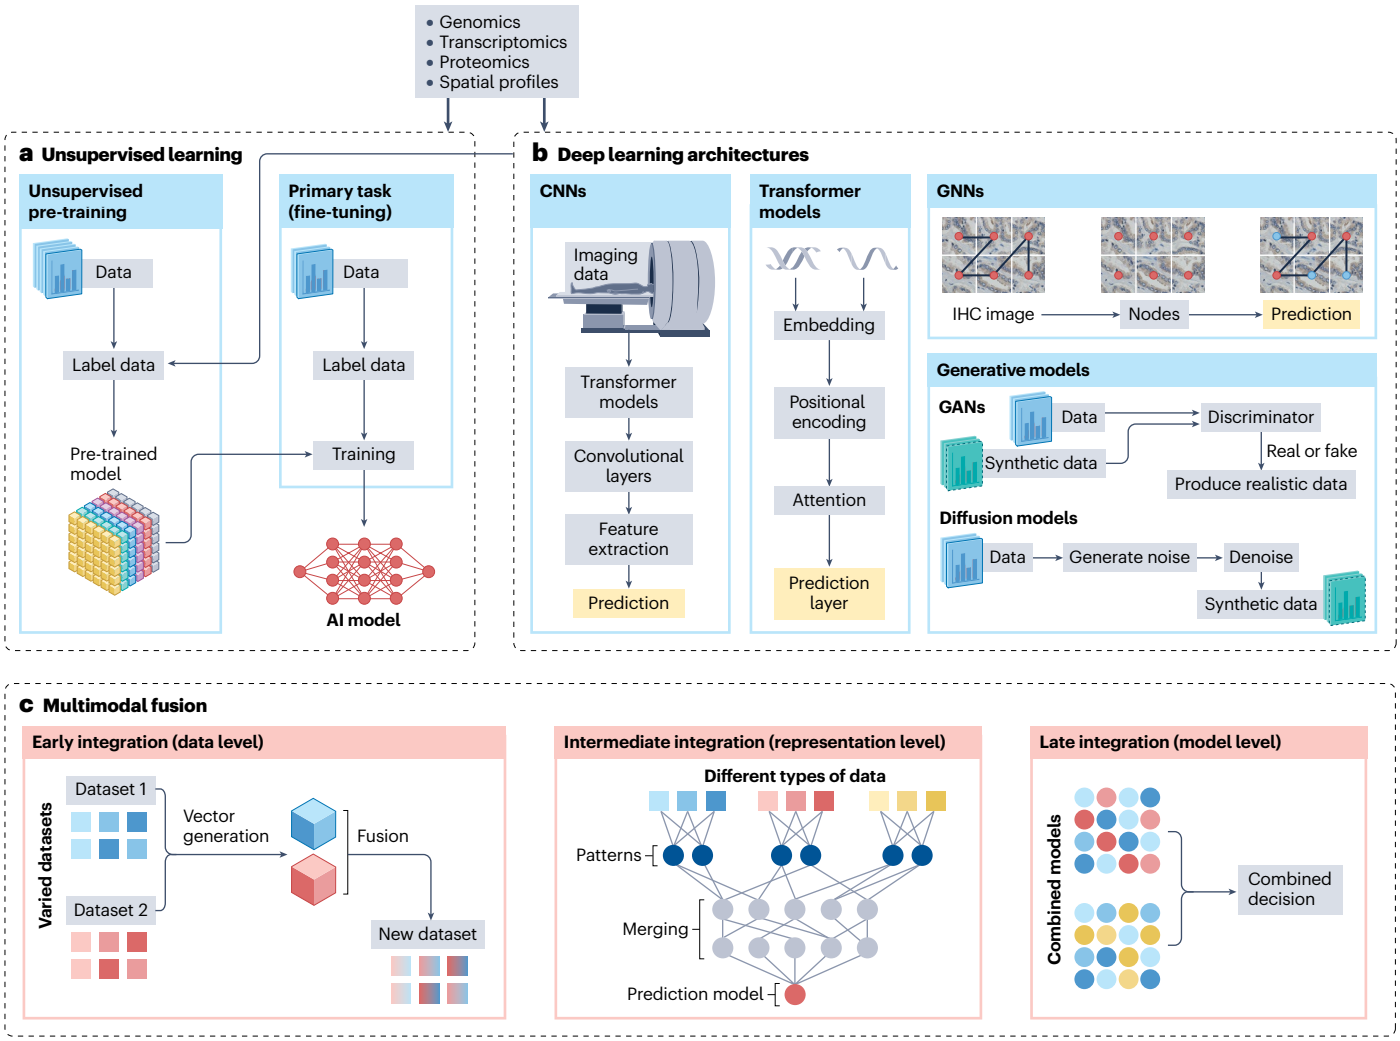

Fig. 2 | AI methodologies for multi-omics data integration. The successful integration of diverse data types in oncology relies on a sophisticated toolkit of artificial intelligence (AI) methodologies. This involves three key methodological dimensions: learning paradigms, model architectures and data fusion strategies. **a**, To address the challenge of data with limited labels, unsupervised or self-supervised pre-training is often employed. Large models are first trained on vast amounts of unlabelled data for an auxiliary task (for example, reconstructing masked inputs), allowing them to learn meaningful data representations. These pre-trained models are then fine-tuned on smaller, labelled datasets for a specific primary task (for example, classification or prediction), generating a foundational AI model that can make new predictions on patient data. **b**, The model architecture defines how the model learns representations during the auxiliary task. A variety of deep learning architectures are available and are selected based on the data modality. Convolutional neural networks (CNNs) are standard for grid-like data such as pathology and radiology images. Transformer and sequence models have a core attention mechanism that enables the dependencies of models across positions in a sequence and are highly effective for sequential data like genomics and transcriptomics. Graph neural networks (GNNs) are uniquely suited to model relational data structures such as biological networks or spatial cell interactions. Generative models, including generative adversarial networks (GANs) and diffusion models, can create realistic synthetic data to augment limited datasets or impute missing modalities. **c**, Multimodal fusion can occur during pre-training or fine-tuning as a part of the model architecture design. Three principal strategies are used to combine information from different sources (for example, genomic data and imaging features). Early integration (data-level) involves extracting and concatenating features into vectors before it is fed into the model. Intermediate integration (representation-level) processes the data from different modalities first through its own encoder to learn patterns specific to that data type. The outputs are then fused in a shared space. Late integration (model-level) involves processing each data type on completely separate models and combining these predictions. The optimal combination of these methodologies depends on the specific biological question and the characteristics of the available data. IHC, immunohistochemistry.

AI model is its architecture, which must be carefully selected to match the characteristics of the input data (Fig. 2b).

Convolutional neural networks (CNNs)[48] have traditionally been the preferred method for extracting local features from grid-like data such as medical images. However, CNNs are generally limited in their ability to model global context, and thus the field is rapidly shifting towards vision transformers (ViTs) and hybrid architectures. While transformers were originally designed for sequential data like

genomics[49], their core 'attention mechanism' is now pivotal for imaging. Unlike CNNs, which focus on local receptive fields and gradually build global context, ViTs divide an image into patches and treat it as a sequence, allowing them to capture long-range dependencies across the entire input. This is particularly advantageous in digital pathology, where determining a prognosis often requires understanding the spatial relationships between tumour cells and distant stromal or immune features across a large (gigapixel) whole-slide image (WSI). Similarly, in 3D radiology, these architectures can better model the volumetric context across slices[50]. Their 'attention mechanism' allows them to weigh the importance of different elements in a long sequence, making them exceptionally skilled at capturing the contextual relationships that define biological function[51]. Conceptually, this mechanism acts as a 'computational spotlight', allowing the model to dynamically focus on the most relevant parts of the input. For example, when analysing a long gene sequence to predict its expression, the attention mechanism can learn to 'pay more attention' to critical regions like promoters or enhancers, rather than treating every base pair with equal importance[52]. This ability to learn long-range, context-dependent relationships is what makes transformers so powerful for biological sequence analysis.

Graph neural networks (GNNs) are uniquely suited to model data with inherent relational or network structures, as they leverage graphs to capture complex dependencies and interactions between the various entities[53]. This contrasts with the grid-based processing of CNNs and the sequential focus of standard transformers. In oncology, biological entities rarely exist in isolation — genes interact in regulatory pathways and proteins form complexes. GNNs can explicitly model these relationships by representing biological entities as nodes and their interactions as edges. This architecture is increasingly applied to integrate multi-omics data onto protein–protein interaction (PPI) networks for cancer driver gene discovery[29] and is becoming the standard for analysing spatial omics data, where it models the tumour microenvironment as a graph of interacting cells to predict patient outcomes[54].

Generative models represent a rapidly emerging frontier. This category includes generative adversarial networks[55] (GANs) and diffusion models[56]. Unlike architectures that primarily predict outcomes from existing data, these models can learn to produce new, high-fidelity synthetic data[57]. The importance of this capability cannot be overstated, particularly for addressing the critical bottleneck of data scarcity in oncology. For many rare cancer subtypes, for instance, datasets may be limited to only a few dozen patients — far too small to train a robust AI model. Generative models can augment these datasets by creating thousands of realistic 'virtual patient' samples, enabling the development of powerful models that would otherwise be impossible. Crucially, their application is broad: these models can generate both highly realistic synthetic pathology images and biologically plausible genomic or transcriptomic sequences, making them a versatile tool across the multi-omics landscape.

Multi-omics fusion

With specialized architectures processing each data type, the final and most critical step is to integrate their insights[58–61]. The choice of fusion strategy determines how the model learns from the interplay between different data streams[62] (Fig. 2c), and this choice involves substantial trade-offs between simplicity, robustness and the ability to uncover complex biological relationships.

Early integration (also termed data-level) is the most direct method, involving the concatenation of raw or processed features from all modalities into a single, wide vector before model input. For instance, a model might combine a vector of 100 radiomic features from a computed tomography (CT) scan with a vector of 20,000 gene-expression values from RNA sequencing (RNAseq) data. While this approach has been used in some early studies[63], its primary drawback is the risk of being dominated by high-dimensional modalities; the 20,000 gene features might computationally 'drown out' the signal from the 100 radiomic features, which makes it difficult for the model to learn subtle but important cross-modal relationships.

Late integration (also termed model-level) is an ensemble technique that takes the opposite approach[37]. It trains a separate, independent model for each modality and then combines their final predictions. For example, a CNN trained on pathology images might predict a 70% chance of recurrence (which can be thought of as its 'vote'), while a separate model trained on genomic data predicts a 30% chance (representing its 'vote'). The final prediction could be a simple majority vote or a weighted average of these outputs. This method is robust and easy to implement, but its major limitation is that it may miss subtle, synergistic effects that are only visible when modalities are considered jointly. For instance, a subtle gene-expression pattern might only be prognostically significant in the presence of a specific morphological feature in the tumour microenvironment. Because late integration analyses each modality in isolation, it is blind to these critical, context-dependent interactions.

Intermediate integration (also termed representation-level integration) has become the most powerful and widely adopted strategy in deep learning precisely because it is designed to find these synergistic effects. Here, each modality first passes through its own specialized encoder — a neural network designed to act as an 'expert summarizer' — to be transformed into a compact, learned representation called an embedding[64]. Conceptually, the encoder for pathology images might be a CNN that summarizes a WSI into an embedding that captures the degree of immune infiltration, while a transformer encoder for RNA sequencing data might create an embedding that summarizes the activity of key immune-related pathways. These separate, expert 'summaries' (the embeddings) are then merged within a shared 'conference room' in the model, known as a hidden layer. Here, the model learns the complex, nonlinear interactions between them. A powerful example is in predicting immunotherapy response, where a model using intermediate integration can learn that the combination of high immune cell density (from the pathology embedding) and a strong interferon-γ gene signature (from the transcriptomic embedding) is a far more powerful predictor of response than either feature alone[19,65].

Finally, the choice of fusion strategy is not about one being definitively superior but about selecting the right tool for the task. Late integration can be a valuable and robust baseline, especially when data modalities are expected to be largely independent. Early integration is rarely used for complex multi-omics data due to its inherent limitations, although it remains a valid and effective approach for integrating lower-dimensional datasets such as combining clinical demographics with targeted biomarker panels. However, intermediate integration is the preferred method when the central scientific goal is to uncover the unknown, synergistic relationships between different layers of cancer biology, which is the core promise of multi-omics research.

Nevertheless, it is crucial to recognize that multimodal integration is not inherently superior to unimodal approaches. The success of fusion depends heavily on the 'orthogonality' of the data — that is, whether the different modalities provide complementary biological information. If modalities are highly redundant (mutual information is high) or if one modality contributes substantially more noise

than signal, fusion can paradoxically degrade model performance or increase complexity without clinical benefit. For instance, seminal reviews have highlighted cases where adding transcriptomic or clinical data to robust genomic models failed to improve prognostic stratification because the additional layers increased dimensionality and noise rather than the actionable signal[7,37]. Specifically, evidence suggests that, across multiple cancer types, clinical variables can outperform genomic and proteomic features for survival prediction, indicating that the additional dimensionality introduced by omics may sometimes dilute the actionable signal inherent in standard clinical records[66,67]. Therefore, in scenarios with limited sample sizes or high data heterogeneity, a well-tuned unimodal model may be preferable to a complex multimodal architecture.

AI-driven multi-omics insights for cancer care

The integration of multi-omics data through the lens of AI is fundamentally reshaping our approach to critical challenges across the cancer care continuum (Table 1).

Identifying underlying molecular and genetic drivers

AI-driven multi-omics integration is beginning to provide unprecedented insights into the molecular and genetic drivers of cancer[45,68–73]. Uncovering this has been a primary challenge in oncology, as these events are often hidden within vast, high-dimensional datasets. First, AI can be used to identify the functional 'footprint' of driver events directly from genomic data, even when that data is complex and weakly labelled. For instance, weakly supervised multiple-instance learning (MIL) models organize data into 'bags' that represent sets of instances, where labels are provided only at the patient ('bag') level and not for individual mutations ('instances'). These models can be trained on the entire 'bag' of somatic mutations from a tumour sample to predict sample-level characteristics, like tumour type or microsatellite status, with high accuracy[74]. This demonstrates how AI can learn to identify the collective mutational patterns that define a cancer's identity without needing a label for each individual mutation. Moving beyond just the presence of mutations, unsupervised deep learning can also decipher their functional impact. In a large preclinical study, the DeepProfile framework was applied to over 50,000 transcriptomes over 18 cancer types, demonstrating that specific gene-expression programmes, represented as latent variables, were consistently associated with both mutation burden and patient survival[20]. This highlights AI's ability to connect the underlying genomic driver landscape to its functional transcriptional consequences.

Furthermore, the use of AI is providing novel insights by deciphering the broader context in which driver genes operate, linking molecular alterations to both visual phenotypes and functional networks. On the visual front, the use of AI has uncovered novel connections between a tumour's driver and its morphology. Seminal studies have demonstrated that deep learning models can predict the status of key driver genes (for example, *EGFR* (encoding epidermal growth factor receptor) or *KRAS*) and microsatellite instability directly from standard hematoxylin and eosin (H&E) slides across multiple cancer types[75,76]. Building on these predictive capabilities, generative approaches are now further bridging this gap. For instance, the GigaTIME framework utilizes multimodal AI to generate virtual spatial proteomics (multiplex immunofluorescence) from standard H&E slides, enabling population-scale modelling of the tumour immune microenvironment to uncover novel protein-level drivers of progression and survival[77]. In parallel, the use of AI can also enable elucidation of the functional context of drivers by modelling the tumour as a biological network using graph-based deep learning. For example, the TREE model was designed to identify cancer driver genes in a pan-cancer context by learning from their positions within heterogeneous biological networks and explicitly integrating multi-omics data (including microRNA (miRNA) expression, transcription factor activity and protein profiles)[29]. Notably, the model's attention mechanism identified that specific network paths, which represent sequential chains of biological signalling such as the regulatory cascade between a transcription factor and its target genes, were highly predictive of driver status. Together, these approaches demonstrate that, by capturing the full context — be it visual or relational — AI can uncover driver information that is invisible to context-unaware analyses.

Detection and diagnosis

Artificial intelligence is providing new pathways towards predictive and non-invasive cancer diagnosis, primarily through the deep analysis of circulating biomarkers from liquid biopsies and the inference of molecular information from medical images[44,78–94]. This is because AI enables the identification of subtle, high-dimensional patterns that are otherwise invisible to traditional analysis. For example, Fragle was initially developed to quantify circulating tumour DNA (ctDNA) in colorectal cancer by analysing DNA-fragment length distributions (fragmentomics). The authors further demonstrated its versatility by applying it to samples from patients with resected lung cancer, where it outperformed a standard gene panel in predicting minimal residual disease (MRD)[95]. However, liquid biopsies exhibit variable sensitivities, particularly for stage I cancers where the ctDNA signal is weakest[96,97]. This creates a clear need for multi-omics integration. Specifically, integrating radiomic features from standard-of-care imaging (for example, CT scans), which capture macroscopic tumour phenotypes, with the microscopic specificity of ctDNA can overcome the sensitivity limitations of single-modality assays[98]. Since a negative liquid biopsy result in an early-stage context cannot definitively rule out cancer, this complementary imaging data is essential for a confident diagnosis. Despite these sensitivity hurdles, the underlying AI capabilities are already powering the development of multi-cancer early detection (MCED) tests that could revolutionize population-level screening[99]. Notably, large-scale prospective trials, such as PATHFINDER, have demonstrated that current AI-driven MCED tests can detect diverse cancer types in asymptomatic individuals with high specificity (>99%) and accurately predict the tissue of origin, thereby effectively directing subsequent diagnostic workups[100].

Beyond genomic signals, AI is increasingly used to harness the diagnostic potential of other circulating analytes, such as proteins and rare cells. For example, in a large-scale pan-cancer study, machine learning classifiers were trained on plasma proteomic profiles from 2,251 patients with cancer. By analysing high-dimensional protein expression patterns, the model successfully identified novel biomarker panels capable of distinguishing multiple tumour types from healthy controls, offering a functional readout of the host–tumour interaction[101]. While such proteomic approaches hold promise, further validation in diverse cohorts is ongoing to translate these findings into routine clinical practice. Furthermore, the use of AI has the potential to enhance the detection of rare circulating cells, which has traditionally been labour-intensive[102]. In this context, novel deep learning approaches have been developed to automatically identify and classify circulating tumour cells (CTCs) from high-content imaging data with expert-level accuracy[103]. This capability is critical because it addresses

Table 1 | Representative examples of AI-driven multi-omics applications in oncology

Oncological application	Cancer type	Data modality	AI/ML method used	Key findings and outcomes (research stage)
Molecular and genetic drivers	Pan-cancer	Proteomics, genomics, RNA sequencing	Graph neural network-based deep learning	Mapped functional protein networks across 11 cancer types using proteogenomic data from 1,194 patients (preclinical)[4]
	Pan-cancer	Multi-omics networks (miRNA, transcription factor, proteins)	Transformer-based graph representation learning	Achieved state-of-the-art performance across 8 pan-cancer datasets and identified 57 high-confidence candidate driver genes (preclinical)[29]
	Breast cancer	Histopathology, molecular profiling	Explainable random forest	Combined morphological and molecular features for comprehensive cancer profiling (retrospective clinical study)[8]
Diagnosis	Lung cancer	CT radiomics, cfDNA fragmentomics, clinical features	Multi-omics model (clinic-RadmC)	Achieved an AUC of 0.923 in predicting malignancy in indeterminate pulmonary nodules, enabling a 10–35% reduction in unnecessary invasive procedures across a multi-institutional cohort (validated in real-world, multi-institutional study)[9]
	Breast cancer	Mammography, clinical workflow	AI-supported screen-reading system	In a randomized trial (MASAI), AI reduced screen-reading workload by 44% while maintaining cancer detection rates (prospective clinical study)[90]
	Pan-cancer (including rare cancers)	Histopathology, clinical metadata	Foundation model (Virchow)	Enabled pan-cancer detection (AUC 0.95) and outperformed tissue-specific models for some rare cancer variants (real-world clinical study)[44]
Precision risk stratification	High-grade serous ovarian cancer	Histopathology, CT imaging, genomics	Multimodal ensemble machine learning	Outperformed single modalities for risk stratification (real-world clinical study)[14]
	Colorectal cancer	Histopathology, multi-omics profiles	Deep learning platform (MOMA)	Predicted survival, microsatellite instability and genomic alterations (real-world clinical study)[13]
	Pan-cancer	Histopathology, genomics	Multimodal deep learning with attention	Achieved effective patient risk stratification across 14 cancer types (c-Index 0.645) and identified interpretable morphological and molecular correlates of survival (real-world clinical study)[17]
	Breast cancer	Mammography, ultrasound, clinical	Multimodal machine learning model	Performed similarly to radiologists for malignancy classification and was superior for differential diagnosis (prospective validation trials)[178]
Personalizing cancer treatment	NSCLC	CT radiomics, PDL1 IHC, genomics (TMB)	Expert-guided ensemble with intermediate integration	Achieved an AUC of 0.80 for immunotherapy response prediction, outperforming either TMB or PDL1 used alone (real-world clinical study)[19]
	Pan-cancer	EHR, genomics, radiology reports	NLP and explainable AI with gradient boosting	Identified *SETD2* mutation as a predictor of immunotherapy response in lung cancer and improved overall survival prediction across 5 cancer types using real-world data[18]
	Pan-cancer (spinal metastases)	CT imaging, radiation dose plans, clinical documentation	Deep learning-based quality assurance (DL-SpiQA)	Prospectively deployed across a multicentre network, the system reviewed treatments in real-time and successfully identified 6 documentation errors and verified targeting accuracy in 520 spine radiation treatments (real-world clinical implementation)[179]
	Prostate cancer	Multi-omics, clinical data, imaging	Multimodal deep learning	Personalized therapy selection in patients based on integrated clinical trial data (validation in phase III randomized controlled trial)[132]
Enhancing prognostic prediction	Endometrial cancer	Histopathology WSI, clinical stage, molecular profiling	Multimodal deep CNN (HECTOR)	Combined pathology and clinical data achieving C-index 0.79–0.83, outperforming molecular profiling alone (validated in real-world randomized trials)[137]
	Pancreatic cancer	Transcriptomics, proteomics, clinical data	Ensemble deep neural networks (Molecular Twin)	Achieved superior disease survival prediction accuracy compared to single-omics or clinical baselines and validated a streamlined 589-feature panel for clinical translation (real-world clinical study)[12]
	NSCLC	Histopathology, CT imaging	Multimodal CNN fusion	Improved survival prediction over the single modalities (real-world clinical study)[39]
	Pan-cancer	EHR, proteomics	Transfer learning (COMET framework)	Achieved an AUROC of 0.842 for predicting 3-year cancer mortality (real-world clinical study)[136]

This table provides representative examples selected to highlight a diverse range of artificial intelligence (AI)-driven applications that explicitly integrate multi-omics data in oncology. The focus is on recent, high-impact studies that demonstrate the synergy between different data types. The 'Stage' refers to the developmental or validation stage of the reported study. AUC, area under the curve; AUROC, area under the receiver operating characteristic curve; cfDNA, circulating cell-free DNA; CNN, convolutional neural network; CT, computed tomography; EHR, electronic health record; IHC, immunohistochemistry; miRNA, microRNA; ML, machine learning; NLP, natural language processing; NSCLC, non-small-cell lung cancer; PDL1, programmed cell death 1 ligand 1; PPI, protein–protein interaction; TMB, tumour mutational burden; WSI, whole-slide image.

the challenge of morphological heterogeneity, allowing for the detection of rare cells that might be missed by manual observation. Capturing these cells provides definitive cytological evidence of tumour dissemination — representing the 'seeds' of metastasis — thereby offering a complementary layer of risk assessment to cell-free DNA assays.

Simultaneously, the use of AI is transforming diagnostic imaging by enabling the inference of the molecular state of a tumour directly from its visual phenotype. This was demonstrated in a large, multicentre retrospective study of patients with lung cancer using DeepGEM, an annotation-free multiple-instance learning model[33]. This model predicted the mutation status of key driver genes (for example, *EGFR* and *KRAS*) from routine H&E-stained pathology slides with high accuracy, achieving an area under the curve (AUC) ranging from 0.90 to 0.97. However, a key challenge with such image-only models is the risk of them learning non-causal 'shortcut' patterns that do not generalize well or are not truly linked to the underlying biology. This is a critical area where multimodal integration is needed; by training a model on both pathology images and the underlying genomic data simultaneously, the genomic data can act as a 'ground truth' to regularize the model and ensure it learns biologically plausible morphological features. Similarly, in radiology, AI algorithms can be used to analyse radiomic features from CT scans and have been successfully used in lung cancer screening programmes to accurately differentiate between benign and malignant lesions[104–109].

Perhaps the synergistic potential of integrating these modalities is the ultimate goal, where the limitations of one modality can be overcome by the strengths of another. This has been exemplified in the clinical challenge of diagnosing indeterminate pulmonary nodules. In a multi-institutional study, clinic-RadmC, an integrated multi-omics model, was developed that combines clinical data, deep learning-based radiomic features from CT scans, and circulating cell-free DNA fragmentomic features in 5-methylcytosine (5mC)-enriched regions[9]. This integrated model achieved a substantially higher diagnostic accuracy (AUC of 0.923) in distinguishing malignant from benign lung nodules compared to any single-omics model. This validated its utility for the accurate diagnosis of lung cancer in patients presenting with indeterminate pulmonary nodules.

Precision risk stratification

Effective cancer management hinges on accurate patient stratification, and while the anatomical tumour node metastasis (TNM) staging system remains essential for this, it does not account for the underlying tumour biology, resulting in the grouping of patients with cancers of distinct molecular profiles and divergent clinical outcomes[110]. This well-documented intra-stage heterogeneity can lead to challenges in personalizing adjuvant therapy, potentially resulting in over-treatment for some patients and under-treatment for others[111]. Molecular subtyping based on single-omics data, such as transcriptomic profiles, has been a major advance in addressing this challenge. However, these approaches do not capture the full complexity of a tumour. Recent frameworks have begun to address this by leveraging AI to integrate disparate clinical and pathological factors for precise risk stratification, as demonstrated in the perioperative assessment of potentially malignant oral disorders to guide surgical interventions[112]. Nevertheless, substantial opportunities remain to deepen this integration by explicitly incorporating multi-omics data, thereby providing a more granular and biologically informed view of patient risk.

These integrative frameworks move decisively beyond both anatomical staging and single-omics subtyping by integrating molecular data with pathomic and radiomic features. For instance, studies have successfully applied this approach to predict therapeutic response in gastric cancer, demonstrating clear superiority over clinical baselines[113]. The power of this approach is amplified when AI extracts thousands of quantitative morphological features, such as features related to nuclear shape, tissue texture and spatial arrangement of cells from WSIs, that serve as powerful surrogates for the underlying tumour biology[114]. A compelling realization of this synergy is seen in high-grade serous ovarian cancer, where a multimodal machine learning model integrating histopathology (WSIs), radiology (CT scans) and genomic features was shown to improve risk stratification compared to models based on any single modality alone[14]. Such integrative strategies are particularly urgent for malignancies like cervical or bladder cancer, where treatment decisions remain heavily tethered to traditional histology. By adapting the multimodal architectures validated in ovarian and gastric models to these contexts, we can unlock a new dimension of diagnostic precision that anatomical staging alone cannot provide.

Enhanced precision in risk stratification can enable the discovery of novel patient subgroups, with unique clinical trajectories and potential therapeutic vulnerabilities[115,116]. For clinicians, this translates into more personalized treatment decisions, such as identifying patients at high risk who may benefit from adjuvant therapy or, conversely, patients at low risk for whom treatment de-escalation could be safely considered, thereby sparing them unnecessary toxicity[117]. Furthermore, in the long term, this approach may also prove more economical and lead to smarter clinical trial designs. For instance, in resource-limited settings, AI models that can infer molecular risk from standard H&E slides could serve as a cost-effective alternative for risk stratification to expensive genomic assays.

Personalized cancer treatment

AI is providing powerful new models to predict patient responses to complex therapeutic regimens, particularly immune-checkpoint inhibitors (ICIs) and targeted therapies, thereby enabling a more precise and individualized approach to therapeutic selection[118–124]. A paramount application lies in predicting response to ICIs, where current biomarkers like programmed cell death 1 ligand 1 (PDL1) expression and tumour mutational burden (TMB) are imperfect predictors. In a large retrospective study of patients with advanced non-small-cell lung cancer (NSCLC), a multimodal model was developed using an intermediate integration strategy to fuse features from radiology (CT scans), pathology (IHC of PDL1 expression), and genomics[19]. This integrated model drastically outperformed unimodal measures, including TMB (AUC 0.61) and PDL1 score (AUC 0.73), in predicting patient response (AUC 0.80). This highlights how fusing macroscopic tumour features from radiology with microscopic and molecular data can serve as a robust decision-support tool to accurately stratify responders, thereby optimizing therapeutic selection. However, prospective validation is still needed before this can be translated to the clinic.

Similarly, AI can analyse tumour transcriptomes to generate gene-expression signatures that predict sensitivity to individual chemotherapy regimens. For example, in a study utilizing preclinical models of pancreatic ductal adenocarcinoma (PDAC), AI-assisted transcriptomic signatures were developed to predict response to gemcitabine and a modified FOLFIRINOX regimen (folinic acid, 5-fluorouracil, irinotecan and oxaliplatin; mFFX)[125]. These signatures were then validated in a phase III clinical trial of 343 patients with PDAC, demonstrating that the patients predicted to be sensitive to the administered drug had a dramatically longer disease-free survival

(median of 50.0 months for mFFX-sensitive patients treated with mFFX versus 10.6 months for patients treated with a regimen to which their tumour was not predicted to be sensitive). This demonstrates a direct path to clinical translation, offering a tool to personalize adjuvant chemotherapy and spare patients from the toxicity of ineffective treatments. This strategy represents a methodological advancement by employing machine learning (specifically LASSO-random forest algorithms) to deconvolute complex transcriptomic data, integrating signals from both neoplastic cells and the tumour microenvironment. By navigating the high-dimensional noise that often limits traditional biomarkers, this AI-driven framework successfully isolates robust, drug-specific predictive signatures essential for multidrug regimens like FOLFIRINOX or mFFX.

Beyond predicting response to existing drugs, the use of AI is also enabling novel, rational therapeutic strategies as biologically informed deep learning models can reveal mechanisms of resistance to targeted therapies[118,126]. A visible neural network model, NeST-VNN, was applied to genomic data from various breast cancer cell lines to generate integrated multi-protein assembly maps and elucidate resistance mechanisms to CDK4 and CDK6 inhibitors[127]. The model identified eight specific core assemblies, including a histone regulatory complex involving KAT6A and TBL1XR1, that mediated resistance. The deployment of AI was essential here because resistance to CDK4 and CDK6 inhibitors is not driven by a single gene but by a constellation of rare and common mutations across dozens of genes. While traditional single-gene biomarkers, such as *RB1* mutation or *CCND1* amplification, failed to translate to clinical populations, this AI-driven approach successfully integrated these scattered genetic signals into functional assembly states to stratify patient outcomes. This approach, which goes beyond single-gene biomarkers, provides specific, testable hypotheses about resistance pathways and could lead to new suggestions of rational therapy combinations that could be investigated in further studies.

Building predictive architectures for cancer prognosis

While prognostication — the prediction of a cancer's likely course — has mostly been based on statistical averages, the use of AI is enabling a transition to a precise, data-driven science[128]. This has been demonstrated in multiple cancer types, including colorectal, breast and prostate cancer, where AI-based models have successfully forecasted individual patient trajectories with high accuracy (AUCs typically ranging from 0.70 to 0.94)[129–133]. This superior performance stems from the ability of CNNs to distil 'deep prognostic' signals (patterns often imperceptible to human experts) from vast histological datasets. Specifically, these algorithms have identified latent morphological biomarkers linked to survival such as tumour-associated adipose features[129], senescence-associated nuclear phenotypes[130], and the complex spatial organization of immune cells within the tumour microenvironment[133]. By capturing biological information inaccessible to standard staging or grading, such pathology-based tools provide independent risk assessments that notably refine clinical decision-making[134,135]. Beyond pathology, similar computational architectures are proving equally effective in chest radiology, where radiomic features were able to serve as non-invasive prognostic biomarkers[39].

The most powerful predictive architectures, however, are those that integrate multiple data types. By fusing molecular data (for example, genomic alterations) with phenotypic data (for example, pathological and radiomic features) and clinical variables, AI can construct a holistic view of the biology of the tumour[136,137]. This integrated approach is particularly valuable in highly heterogeneous cancers where a single data modality provides an incomplete picture. For example, in a study of resected PDAC[12], a precision medicine platform known as the 'Molecular Twin' used machine learning to integrate thousands of multi-omics features (including proteomics) with clinical data. This analysis revealed that, while plasma protein profiles emerged as the superior single-omics predictor, the integrated multimodal model achieved the highest prognostic accuracy, drastically outperforming any individual data source. Critically, these architectures are not static, meaning that they lay the groundwork for dynamic prognostic tools that can be updated over time with new longitudinal data, such as changes in ctDNA levels, allowing for the continuous reassessment of a patient's risk and a more adaptive approach to their long-term management[138].

Proactive surveillance

Historically, post-treatment surveillance has been a reactive process, relying on scheduled imaging to detect macroscopic recurrence, often when the window for curative-intent intervention has narrowed[139]. AI-driven analysis of longitudinal molecular data is ushering in a new era of proactive surveillance[140,141]. The primary computational challenge in this domain is distinguishing infinitesimally small tumour-derived signals from a noisy background of normal cell-free DNA. Sophisticated deep learning algorithms have been purpose-built to solve this signal-to-noise problem, analysing fragmentomic patterns to track ctDNA levels with a sensitivity that far surpasses conventional imaging[95].

Its clinical implications are profound. Firstly, longitudinal ctDNA monitoring allows for the identification of impending relapse months before it becomes radiologically evident, creating an opportunity for early therapeutic intervention when the disease burden is lowest[142]. Secondly, interpretable AI models integrating liquid biopsy with clinical features can refine progression risk prediction and identify patients most likely to benefit from adjuvant therapy[143]. However, liquid biopsy alone lacks spatial information regarding the tumour microenvironment. This is where multimodal AI integration proves transformative. A compelling example is the 'HIBRID' framework developed for colorectal cancer, which synergizes deep learning-based risk scores derived from routine histology (H&E slides) with ctDNA-based MRD status[144]. Uniquely, this integrated architecture demonstrated that patients who were MRD negative but classified as being at high risk by the deep learning histology model still derived noteworthy benefit from adjuvant chemotherapy — a critical subgroup that would have been overlooked by ctDNA monitoring alone. This shift from static, single-source observation to dynamic, AI-unified, multimodal monitoring represents a crucial step towards the long-term management of cancer.

The oncological digital twin for personalized prediction

The ultimate synthesis of AI-driven multi-omics integration is the creation of a dynamic, computational replica that mirrors an individual's cancer biology, known as a 'digital twin'[145–149]. This concept, originating from engineering where virtual models of complex systems like jet engines are used for in silico testing[150], is now being adapted for oncology[151]. It is not merely a data repository but a functional, predictive model that represents the next frontier in personalized medicine. A foundational step towards this goal has been demonstrated by models like LifeClock[31], which integrated millions of heterogeneous, longitudinal electronic health records (EHRs) to build a universal biological clock capable of predicting health risks across the entire human lifespan. This illustrates the core principle of a digital twin, which is to generate

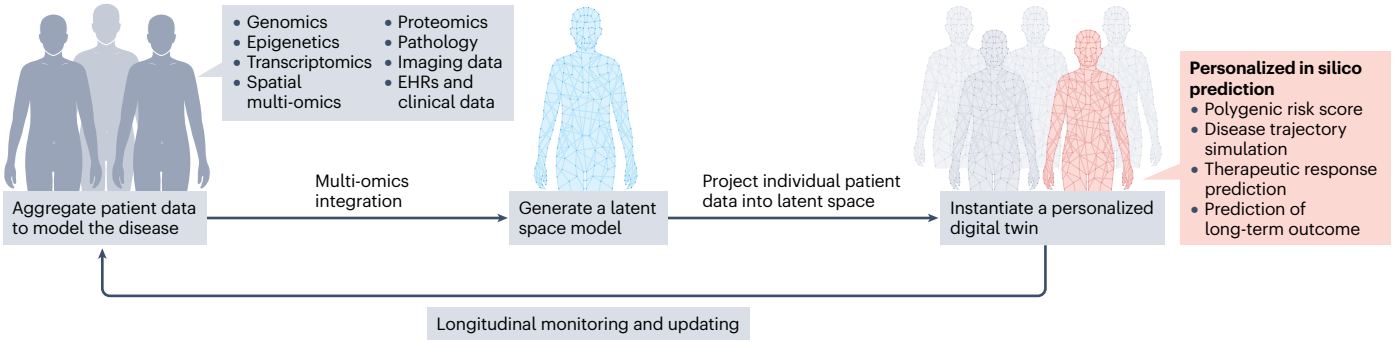

Fig. 3 | The conceptual framework of developing and utilizing digital twins in oncology. The creation of a patient digital twin is a dynamic process that translates multi-scale biology into a functional computational model. This process begins by aggregating large cohorts of diverse patient data — spanning molecular, clinical and imaging modalities — to train a latent space model. This latent space acts as a compressed, lower-dimensional representation of high-dimensional data that captures the most essential patterns and relationships within the data. Therefore, this serves as a conceptual 'map' of the disease landscape for the artificial intelligence (AI) model. To instantiate a personalized digital twin, an individual patient's available data is projected into this pre-trained latent space. Crucially, this deep learning architecture is capable of handling incomplete datasets (for example, if only genomic data and basic demographics are available); it uses the learned population-level structure to infer missing modalities, thereby generating a comprehensive deep phenotype from limited input data. Once instantiated, the digital twin serves as a dynamic replica for personalized in silico prediction. Rather than simply categorizing risk, the model functions as a simulation engine, allowing clinicians to model future disease trajectories and test therapeutic responses computationally. The entire framework operates within a continuous feedback loop of longitudinal monitoring and updating, where new data from the patient's ongoing journey is used to constantly recalibrate their digital twin, ensuring it remains an accurate and evolving representation of their biological status. EHRs, electronic health records.

a general model of a biological process by analysing data from a large cohort, which can then be applied to accurately describe the unique characteristics of a single individual[152].

The construction and utility of such a twin can be understood as a three-stage process (Fig. 3): first, the model learns a general representation of the disease from large cohorts; next, this knowledge is used to create a personalized model for an individual patient; and, finally, the personalized model is deployed for in silico simulation and prediction. The foundation of any robust digital twin is the collective knowledge learned from a large and diverse patient population[153]. The process begins with the aggregation of high-dimensional, multi-omics data from thousands of patients into a digital databank. Deep learning models are then trained on this entire cohort to distil its immense complexity into a structured, meaningful representation known as the latent space[154]. This can be conceptualized as a compressed, multi-dimensional 'map' of the disease, capturing the fundamental rules and relationships that govern cancer biology.

Next, deep phenotyping is used to generate an individual's digital twin[155,156]. This involves taking the patient's comprehensive multi-omics data, often encompassing clinical, genomic, imaging and histopathological information[157], and projecting it into the pre-existing latent space. Crucially, this projection does not require perfect data completeness; the model leverages the learned population-level structure to infer missing modalities from available inputs (for example, predicting molecular states from histology), enabling robust twin instantiation even with partial datasets. The AI model then embeds the patient within the 'map', assigning them a unique coordinate. This coordinate is the patient's deep phenotype, which should ideally capture complex patient-specific states such as the tumour's immune activation level or its metabolic profile. This is the critical instantiation step, creating a concrete and computationally tractable model of a single individual that is ready for computational interrogation.

The true power of the instantiated digital twin lies in its utility as a dynamic, predictive tool. Because the twin is a functional model, it can be perturbed and manipulated in silico to simulate future events and guide clinical decisions[158,159]. This enables a form of personalized in silico clinical trial, where clinicians can test the potential outcomes of multiple therapeutic strategies. For example, mechanistic models, which simulate the underlying biological processes of tumour growth and drug response, are a key component of many digital twins[160]. In a powerful preclinical demonstration for high-risk paediatric liver tumours, a digital twin was created for each patient by integrating their clinical, genetic and transcriptomic data. By running in silico drug-response simulations on these patient-specific models, the system identified ceritinib as a novel and effective treatment option, a finding that was subsequently validated in patient-derived xenograft mouse models[161]. This illustrates how digital twins can not only predict responses to standard therapies but can also discover non-obvious drug candidates. This capability allows for the direct comparison of different treatment options in silico, identifying the optimal strategy for an individual patient. This could ultimately move oncology from a reactive discipline to a proactive and predictive one, fulfilling the ultimate promise of truly personalized cancer medicine.

The role of XAI

Complex AI models have undeniable predictive power in multi-omics analysis; however, their clinical translation is often hampered by their inherent lack of explainability, or 'black box' nature, as trust and safety are of paramount importance. This opacity arises from the same non-linear complexity that makes the models so effective[162,163]. Clinicians and regulators need confidence that a model's predictions are based on biologically relevant features, not on spurious correlations or technical artifacts[28]. XAI provides a suite of techniques designed to demystify these models[164]. A compelling example is the application of XAI to a DNA methylation-based classifier for brain tumours. Here, an interpretable framework was used to dissect the model's decision process, revealing that it relied on functionally relevant genomic regions, such as enhancers and large-scale heterochromatic domains, to distinguish

between tumour classes[165]. By confirming that the model's logic aligns with established principles of epigenomic regulation, this XAI analysis directly built the trust necessary for its responsible integration into clinical workflows[166].

However, the role of XAI extends beyond mere validation. It can transform predictive models into powerful engines for biological discovery[165]. A highly accurate model may identify unexpected cross-modal features as critically important for its predictions, generating novel, data-driven hypotheses[167]. For example, biologically informed neural networks were used to predict prostate cancer aggressiveness by integrating gene-expression data with a known hierarchy of biological pathways. This generated an interpretable deep learning model that was able to not only predict patient outcomes but also to identify novel molecular drivers of treatment resistance, such as MDM4 and FGFR1, which were subsequently validated in vitro[168]. This demonstrates that XAI does not simply return a list of important genes from one modality. Instead, it explains how the model integrates information across biological scales — from gene expression to pathway activation — to make its prediction. This capability suggests that AI can be more than a simple predictive tool, but rather a collaborative affiliate in science, allowing researchers to learn directly from the model and discover new, systems-level insights that can be used to improve our understanding of cancer biology[169,170].

Perspectives and conclusion

The integration of AI with multi-omics is not merely creating better predictive models; it is fundamentally reshaping the paradigm of cancer research itself[171]. The future lies in creating a dynamic, self-improving discovery ecosystem or an AI-driven oncologic multi-omics loop (Fig. 4). This framework moves beyond the traditional linear path of bench-to-bedside research towards a continuous, cyclical learning process[172]. It begins with AI models discovering novel patterns and associations from real-world clinical and molecular data, which in turn allows researchers to generate precise, data-driven biological hypotheses. These hypotheses can then be rigorously tested in clinical trials, generating new, high-quality multi-omics data[173]. This new data feeds back into the loop, where AI is again employed for deep mechanistic elucidation and hypothesis validation. The resulting knowledge, now clinically and biologically validated, is integrated back into clinical practice, enriching the data pool and initiating the next cycle of discovery. We believe that this iterative loop will accelerate the pace of innovation, ensuring that clinical practice and biological understanding co-evolve in a rapid and robust manner.

Glossary

Convolutional neural networks
(CNNs). A class of deep learning architectures inspired by the human visual cortex, which are exceptionally effective at learning hierarchical patterns from grid-like data such as pathology and radiology images.

Deep learning
A subfield of artificial intelligence (AI) that uses multi-layered artificial neural networks to learn complex patterns and representations directly from large datasets such as identifying features in medical images or genomic sequences.

Digital twins
Dynamic, computational replicas of individual patients that integrate longitudinal multi-omics and clinical data to simulate disease trajectories and predict responses to therapeutic interventions in silico.

Encoders
An AI model component that transforms raw input data into lower-dimensional, structured representations to capture the most relevant features, reduce noise and facilitate integration or downstream tasks.

Explainable AI
(XAI). A suite of AI methods designed to make the reasoning and predictions of complex models (or 'black boxes') understandable to humans, which is critical for building clinical trust and generating novel biological hypotheses.

Foundation models
Large-scale AI models that are pre-trained on vast amounts of broad, unlabelled data (for example, millions of pathology slides) and can be subsequently fine-tuned with minimal task-specific data for a wide range of downstream applications.

Generative models
A class of AI models, including generative adversarial networks and diffusion models, that can create new, high-fidelity synthetic data to augment limited datasets or impute missing modalities.

Graph neural networks
(GNNs). A class of deep learning architectures specifically designed to learn from data structured as a graph or network; in oncology, they are widely used to model non-Euclidean data such as biological interaction networks or the spatial arrangement of cells within the tumour microenvironment.

Latent space
A compressed, lower-dimensional representation of high-dimensional data that captures the most essential patterns and relationships within the data, to serve as a conceptual 'map' of the disease landscape for an AI model.

Liquid biopsy
A minimally invasive diagnostic method for cancer screening, diagnosis and monitoring that analyses biomarkers, such as circulating tumour DNA, in bodily fluids like blood.

Minimal residual disease
(MRD). A state where minute quantities of tumour cells persist after treatment and can eventually lead to recurrence, even when the patient is asymptomatic and shows no visible signs of disease on imaging.

Radiomics
A computational approach that converts medical images (such as CT or MRI scans) into high-dimensional data by extracting a large number of quantitative features, which can then be mined for patterns correlated with clinical or molecular characteristics.

Supervised learning
A fundamental machine learning paradigm where a model learns to make predictions by training on a dataset of examples that have been manually labelled with the correct output or ground truth.

Transformers
A class of deep learning architectures, originally from natural language processing, that use an 'attention mechanism' to weigh the importance of different elements in a sequence, making them highly effective for analysing biological data like genomics and transcriptomics.

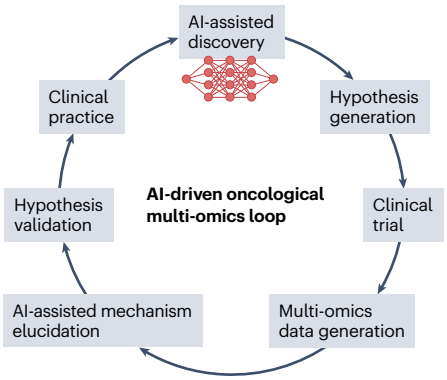

Fig. 4 | The AI-driven oncological multi-omics loop. The process begins with data from routine clinical practice, which serves as a rich source of real-world evidence. Artificial intelligence (AI) models then perform large-scale pattern discovery on these data, identifying novel correlations and signatures that are often imperceptible to human analysis. These data-driven insights enable researchers to formulate a new hypothesis that can then be rigorously tested in a clinical trial. This not only serves to validate the initial idea but also to generate new, high-quality multi-omics data from the trial cohort. In the next phase, AI is deployed again for in-depth mechanism elucidation, analysing the new data to uncover the biological underpinnings of the observed phenomena. This leads to hypothesis validation and the creation of new knowledge, which is then integrated back to inform and improve clinical practice, thus generating a loop and initiating the next cycle of innovation.

Despite this immense potential, the path to widespread clinical implementation is still fraught with substantial hurdles that require a concerted effort from the entire research community. The primary challenge remains rooted in the data: the acquisition of large, high-quality and ethnically diverse multi-omics datasets is essential to train robust and equitable AI models, yet this remains a costly and logistically complex endeavour[174,175]. The issue of model generalizability is also paramount; models trained at one institution often fail to perform well on data from another due to technical or population-based variations[42,47]. Rigorous, prospective clinical validation of AI tools is still largely lacking, and clear regulatory pathways for these complex models are only beginning to emerge. Furthermore, leveraging XAI will be essential for proactively managing these challenges, as it provides the necessary tools to audit models for potential algorithmic bias and to ensure that these powerful technologies are deployed safely, ethically and fairly (Box 1).

Beyond safety and ethics, sustainable clinical implementation hinges on resolving granular operational challenges. Foremost among these is interoperability; AI models must seamlessly exchange data with EHRs using standardized protocols. For example, the Fast Healthcare Interoperability Resources (FHIR) standard[176], developed by HL7 International, provides a modern framework for exchanging structured health data across diverse systems, ensuring seamless integration without disrupting clinical workflows. Economically, the high cost of comprehensive multi-omics profiling and the substantial data infrastructure required for storage and computing remain sizable barriers, necessitating rigorous cost-effectiveness analyses to

Box 1 | Navigating the path to safe and equitable clinical AI

The transition of artificial intelligence (AI) from research to routine oncology practice requires addressing critical challenges in ethics, robustness and governance.

Algorithmic bias and fairness

AI models can inadvertently perpetuate or exacerbate health disparities if trained on non-representative data. For instance, models trained predominantly on populations of European ancestry may perform poorly on other racial or ethnic groups due to underlying differences in genetic ancestry or social determinants of health[174]. Addressing this requires proactive data curation to ensure population diversity and the use of algorithmic fairness frameworks to audit models for subgroup performance disparities[38].

Generalizability and domain adaptation

A major hurdle is 'domain shift', where a model fails when applied to data from a new hospital due to technical variations (for example, different scanners or staining protocols). Domain adaptation techniques allow models to adjust to these variations without retraining from scratch. Furthermore, rigorous external validation across multi-institutional cohorts is non-negotiable to prove that a model learns robust biological signals rather than site-specific artefacts[173].

Model auditing and uncertainty quantification

Trust is built on transparency. Model auditing involves probing AI systems to identify failure modes and 'shortcuts' (non-causal correlations) using generative counterfactuals or expert reviews[162]. Additionally, deploying models with uncertainty quantification allows the AI to express a 'confidence score' for each prediction. In clinical workflows, this enables a 'human-in-the-loop' safeguard: if the model is uncertain about a complex case, it can automatically defer the decision to a clinician rather than making a potentially erroneous guess.

Regulatory and governance frameworks

Safe deployment necessitates standardized evaluation. Evolving guidelines, such as PROBAST+AI, provide structured frameworks to assess the risk of bias and applicability of prediction models[180]. Adhering to such standards is essential for regulatory approval and for ensuring that AI tools are not only accurate but also safe and equitable for all patients.

The lifecycle of responsible AI deployment

To ensure safety and equity, AI development must progress through: (1) curating diverse datasets to represent all patient subgroups; (2) employing domain adaptation to handle technical variability between hospitals; (3) rigorous auditing to detect algorithmic bias before deployment; and (4) integrating uncertainty quantification to defer low-confidence predictions to human clinicians.

justify reimbursement, particularly in the short term, where upfront costs remain prohibitive. Finally, regulatory frameworks face the unique challenge of 'adaptive algorithms' — models that continuously learn and evolve after deployment. Unlike traditional 'locked' medical devices, these dynamic systems will require novel regulatory pathways, such as pre-determined change control plans, to monitor performance drift without requiring full re-approval for every update[177]. While substantial scientific, ethical and operational challenges remain, the convergence of technological innovation and interdisciplinary collaboration is and should continue to pave the way for a future where cancer care is predictive, personalized and proactive. The continued commitment to rigorous science and responsible implementation of AI will be the key to realizing this transformative potential for patients with cancer.

Published online: 21 April 2026

References

1. Yang, Z., Guan, F., Bronk, L. & Zhao, L. Multi-omics approaches for biomarker discovery in predicting the response of esophageal cancer to neoadjuvant therapy: a multidimensional perspective. *Pharmacol. Ther.* **254**, 108591 (2024).
2. The ICGC/TCGA Pan-Cancer Analysis of Whole Genomes Consortium. Pan-cancer analysis of whole genomes. *Nature* **578**, 82–93 (2020).
   This consortium paper provides one of the most comprehensive genomic and transcriptomic datasets across multiple cancers, forming a foundational resource for pan-cancer computational analyses.
3. Charlotte, J. H. & Jeffrey, M. D. Artificial intelligence and machine learning in clinical medicine. *N. Engl. J. Med.* **388**, 1201–1208 (2023).
4. Shi, Z., Lei, J. T., Elizarraras, J. M. & Zhang, B. Mapping the functional network of human cancer through machine learning and pan-cancer proteogenomics. *Nat. Cancer* **6**, 205–222 (2025).
5. Zeng, Q. et al. Understanding tumour endothelial cell heterogeneity and function from single-cell omics. *Nat. Rev. Cancer* **23**, 544–564 (2023).
6. Gomes, B. & Ashley, E. A. Artificial intelligence in molecular medicine. *N. Engl. J. Med.* **388**, 2456–2465 (2023).
7. Lipkova, J. et al. Artificial intelligence for multimodal data integration in oncology. *Cancer Cell* **40**, 1095–1110 (2022).
   This review offers a foundational synopsis of the strategies, applications and challenges of using AI for multimodal data integration in oncology.
8. Binder, A. et al. Morphological and molecular breast cancer profiling through explainable machine learning. *Nat. Mach. Intell.* **3**, 355–366 (2021).
9. Zhao, M. et al. Integrated multiomics signatures to optimize the accurate diagnosis of lung cancer. *Nat. Commun.* **16**, 84 (2025).
   This paper provides a compelling clinical use-case, demonstrating that a multi-omics model integrating radiomics and cell-free DNA fragmentomics improves the diagnostic accuracy for indeterminate pulmonary nodules.
10. Alcazer, V. et al. Evaluation of a machine-learning model based on laboratory parameters for the prediction of acute leukaemia subtypes: a multicentre model development and validation study in France. *Lancet Digit. Health* **6**, e323–e333 (2024).
11. Wang, T. et al. MOGONET integrates multi-omics data using graph convolutional networks allowing patient classification and biomarker identification. *Nat. Commun.* **12**, 3445 (2021).
12. Osipov, A. et al. The Molecular Twin artificial-intelligence platform integrates multi-omic data to predict outcomes for pancreatic adenocarcinoma patients. *Nat. Cancer* **5**, 299–314 (2024).
    This paper introduces the 'Molecular Twin' concept, demonstrating a powerful platform that successfully integrates thousands of multi-omics features to accurately predict survival in pancreatic cancer.
13. Tsai, P. C. et al. Histopathology images predict multi-omics aberrations and prognoses in colorectal cancer patients. *Nat. Commun.* **14**, 2102 (2023).
    A key demonstration of the genotype–phenotype link, this work shows that deep learning on standard histopathology images can predict a wide range of multi-omics aberrations, including gene expression and copy number alterations.
14. Boehm, K. M. et al. Multimodal data integration using machine learning improves risk stratification of high-grade serous ovarian cancer. *Nat. Cancer* **3**, 723–733 (2022).
    This study provides a key example of how fusing histopathological, radiological and clinicogenomic features can improve prognostic accuracy over any single modality alone.
15. Niazi, M. K. K., Parwani, A. V. & Gurcan, M. N. Digital pathology and artificial intelligence. *Lancet Oncol.* **20**, e253–e261 (2019).
16. Sunami, K. et al. A learning program for treatment recommendations by molecular tumor boards and artificial intelligence. *JAMA Oncol.* **10**, 95–102 (2024).
17. Chen, R. J. et al. Pan-cancer integrative histology-genomic analysis via multimodal deep learning. *Cancer Cell* **40**, 865–878.e6 (2022).
18. Jee, J. et al. Automated real-world data integration improves cancer outcome prediction. *Nature* **636**, 728–736 (2024).
    This paper details the creation of a massive, multi-modal real-world dataset (MSK-CHORD) and shows that integrating NLP-derived features from clinical notes improves survival prediction beyond genomics or stage alone.
19. Vanguri, R. S. et al. Multimodal integration of radiology, pathology and genomics for prediction of response to PD-(L)1 blockade in patients with non-small cell lung cancer. *Nat. Cancer* **3**, 1151–1164 (2022).
    This study demonstrated that a model integrating radiology, pathology and genomics can predict immunotherapy response in NSCLC more accurately than established biomarkers like TMB or PDL1.
20. Qiu, W. et al. Deep profiling of gene expression across 18 human cancers. *Nat. Biomed. Eng.* **9**, 333–355 (2025).
21. Chang, T. G., Park, S., Schaffer, A. A., Jiang, P. & Ruppin, E. Hallmarks of artificial intelligence contributions to precision oncology. *Nat. Cancer* **6**, 417–431 (2025).
22. Unger, M. & Kather, J. N. Deep learning in cancer genomics and histopathology. *Genome Med.* **16**, 44 (2024).
23. LeCun, Y., Bengio, Y. & Hinton, G. Deep learning. *Nature* **521**, 436–444 (2015).
24. Bravo González-Blas, C. et al. Single-cell spatial multi-omics and deep learning dissect enhancer-driven gene regulatory networks in liver zonation. *Nat. Cell Biol.* **26**, 153–167 (2024).
25. Wornow, M. et al. Zero-shot clinical trial patient matching with LLMs. *NEJM AI* **2**, AIcs2400360 (2025).
26. Andani, S. et al. Histopathology-based protein multiplex generation using deep learning. *Nat. Mach. Intell.* **7**, 1292–1307 (2025).
27. Kondepudi, A. et al. Foundation models for fast, label-free detection of glioma infiltration. *Nature* **637**, 439–445 (2025).
28. Wang, J. et al. An interpretable artificial intelligence framework for designing synthetic lethality-based anti-cancer combination therapies. *J. Adv. Res.* **65**, 329–343 (2024).
29. Su, X. et al. Interpretable identification of cancer genes across biological networks via transformer-powered graph representation learning. *Nat. Biomed. Eng.* **9**, 371–389 (2025).
    This paper presents a novel approach using a transformer-based model on biological networks to not only predict cancer driver genes with high accuracy but also to provide interpretable insights into their regulatory mechanisms.
30. Prelaj, A. et al. Artificial intelligence for predictive biomarker discovery in immuno-oncology: a systematic review. *Ann. Oncol.* **35**, 29–65 (2024).
31. Wang, K. et al. A full life cycle biological clock based on routine clinical data and its impact in health and diseases. *Nat. Med.* **31**, 4225–4235 (2025).
32. Zhu, J. et al. Cell Decoder: decoding cell identity with multi-scale explainable deep learning. *Genome Biol.* **26**, 359 (2025).
33. Zhao, Y. et al. Deep learning using histological images for gene mutation prediction in lung cancer: a multicentre retrospective study. *Lancet Oncol.* **26**, 136–146 (2025).
    This study demonstrates the clinical potential of predicting key driver gene mutations in lung cancer directly from routine H&E-stained pathology slides with high accuracy across multiple centres.
34. Tan, C. L. et al. Prediction of tumor-reactive T cell receptors from scRNA-seq data for personalized T cell therapy. *Nat. Biotechnol.* **43**, 134–142 (2025).
35. Rakaee, M. et al. Deep learning model for predicting immunotherapy response in advanced non-small cell lung cancer. *JAMA Oncol.* **11**, 109–118 (2025).
36. Ellrott, K. et al. Classification of non-TCGA cancer samples to TCGA molecular subtypes using compact feature sets. *Cancer Cell* **43**, 195–212.e11 (2025).
37. Boehm, K. M., Khosravi, P., Vanguri, R., Gao, J. & Shah, S. P. Harnessing multimodal data integration to advance precision oncology. *Nat. Rev. Cancer* **22**, 114–126 (2022).
38. Chen, R. J. et al. Algorithmic fairness in artificial intelligence for medicine and healthcare. *Nat. Biomed. Eng.* **7**, 719–742 (2023).
39. Nam, J. G. et al. Histopathologic basis for a chest CT deep learning survival prediction model in patients with lung adenocarcinoma. *Radiology* **305**, 441–451 (2022).
40. Cai, Z. et al. Synthetic augmentation of cancer cell line multi-omic datasets using unsupervised deep learning. *Nat. Commun.* **15**, 10390 (2024).
41. Lotfollahi, M., Yuhan, H., Theis, F. J. & Satija, R. The future of rapid and automated single-cell data analysis using reference mapping. *Cell* **187**, 2343–2358 (2024).
42. Wang, H. et al. Scientific discovery in the age of artificial intelligence. *Nature* **620**, 47–60 (2023).
43. Figueiredo, M. A. T. & Jain, A. K. Unsupervised learning of finite mixture models. *IEEE Trans. Pattern Anal. Mach. Intell.* **24**, 381–396 (2002).
44. Vorontsov, E. et al. A foundation model for clinical-grade computational pathology and rare cancers detection. *Nat. Med.* **30**, 2924–2935 (2024).
45. Nguyen, E. et al. Sequence modeling and design from molecular to genome scale with Evo. *Science* **386**, eado9336 (2024).
46. Xu, H. et al. A whole-slide foundation model for digital pathology from real-world data. *Nature* **630**, 181–188 (2024).
47. Fahrner, L. J., Chen, E., Topol, E. & Rajpurkar, P. The generative era of medical AI. *Cell* **188**, 3648–3660 (2025).
48. Krizhevsky, A., Sutskever, I. & Hinton, G. E. ImageNet classification with deep convolutional neural networks. *Adv. Neural Inf. Process. Syst.* **25**, 1097–1105 (2012).
49. Vaswani, A. et al. Attention is all you need. *Adv. Neural Inf. Process. Syst.* **30**, 5998–6008 (2017).

50. Song, A. H. et al. Analysis of 3D pathology samples using weakly supervised AI. *Cell* **187**, 2502–2520.e17 (2024).
51. Szałata, A. et al. Transformers in single-cell omics: a review and new perspectives. *Nat. Methods* **21**, 1430–1443 (2024).
52. Avsec, Ž et al. Effective gene expression prediction from sequence by integrating long-range interactions. *Nat. Methods* **18**, 1196–1203 (2021).
53. Li, H. et al. CGMega: explainable graph neural network framework with attention mechanisms for cancer gene module dissection. *Nat. Commun.* **15**, 5997 (2024).
54. Ritter, M. et al. Spatially resolved transcriptomics and graph-based deep learning improve accuracy of routine CNS tumor diagnostics. *Nat. Cancer* **6**, 292–306 (2025).
55. Goodfellow, I. et al. Generative adversarial networks. *Commun. ACM* **63**, 139–144 (2020).
56. Yang, L. et al. Diffusion models: a comprehensive survey of methods and applications. *ACM Comput. Surv.* **56**, 1–39 (2023).
57. Wang, J. et al. Self-improving generative foundation model for synthetic medical image generation and clinical applications. *Nat. Med.* **31**, 609–617 (2025).
58. Zhao, F., Zhang, C. & Geng, B. Deep multimodal data fusion. *ACM Comput. Surv.* **56**, 1–36 (2024).
59. Prakash, A., Chitta, K. & Geiger, A. in *Proc. IEEE/CVF Conf. Comput. Vision Pattern Recognition* 7077–7087 (IEEE, 2021).
60. Wang, Y. Survey on deep multi-modal data analytics: collaboration, rivalry, and fusion. *ACM Trans. Multimed. Comput. Commun. Appl.* **17**, 1–25 (2021).
61. Yang, H. et al. A multimodal vision–language model for generalizable annotation-free pathology localization. *Nat. Biomed. Eng.* https://doi.org/10.1038/s41551-025-01574-7 (2026).
62. Gadzicki, K., Khamsehashari, R. & Zetzsche, C. in *2020 IEEE 23rd Int. Conf. Information Fusion (FUSION)* 1–6 (IEEE, 2020).
63. Chaudhary, K., Poirion, O. B., Lu, L. & Garmire, L. X. Deep learning-based multi-omics integration robustly predicts survival in liver cancer. *Clin. Cancer Res.* **24**, 1248–1259 (2018).
64. Zhu, J., Huang, C. & De Meo, P. DFMKE: a dual fusion multi-modal knowledge graph embedding framework for entity alignment. *Inf. Fusion.* **90**, 111–119 (2023).
65. Wang, M., Bai, S. & Zhang, Y. Single- and multi-modal molecular probes with second near-infrared activatable optical signals for disease diagnosis and theranostics. *Chem. Soc. Rev.* **54**, 7561–7609 (2025).
66. Herrmann, M., Probst, P., Hornung, R., Jurinovic, V. & Boulesteix, A. L. Large-scale benchmark study of survival prediction methods using multi-omics data. *Brief. Bioinform.* **22**, bbaa167 (2021).
67. Yuan, Y. et al. Assessing the clinical utility of cancer genomic and proteomic data across tumor types. *Nat. Biotechnol.* **32**, 644–652 (2014).
68. Cheng, J. et al. Accurate proteome-wide missense variant effect prediction with AlphaMissense. *Science* **381**, eadg7492 (2023).
69. Carrillo-Perez, F. et al. Generation of synthetic whole-slide image tiles of tumours from RNA-sequencing data via cascaded diffusion models. *Nat. Biomed. Eng.* **9**, 320–332 (2025).
A key generative AI application, this work demonstrates the novel capability of generating realistic histopathology images directly from RNA-sequencing data, offering a powerful solution for data scarcity and modality imputation.
70. Howard, F. M. et al. Generative adversarial networks accurately reconstruct pan-cancer histology from pathologic, genomic, and radiographic latent features. *Sci. Adv.* **10**, eadq0856 (2024).
71. Lorenzo, G. et al. Patient-specific, mechanistic models of tumor growth incorporating artificial intelligence and big data. *Annu. Rev. Biomed. Eng.* **26**, 529–560 (2024).
72. Hu, B. et al. High-resolution spatially resolved proteomics of complex tissues based on microfluidics and transfer learning. *Cell* **188**, 734–748.e22 (2025).
73. He, B. et al. Integrating spatial gene expression and breast tumour morphology via deep learning. *Nat. Biomed. Eng.* **4**, 827–834 (2020).
74. Anaya, J., Sidhom, J. W., Mahmood, F. & Baras, A. S. Multiple-instance learning of somatic mutations for the classification of tumour type and the prediction of microsatellite status. *Nat. Biomed. Eng.* **8**, 57–67 (2024).
75. Coudray, N. et al. Classification and mutation prediction from non-small cell lung cancer histopathology images using deep learning. *Nat. Med.* **24**, 1559–1567 (2018).
76. Kather, J. N. et al. Pan-cancer image-based detection of clinically actionable genetic alterations. *Nat. Cancer* **1**, 789–799 (2020).
77. Valanarasu, J. M. J. et al. Multimodal AI generates virtual population for tumor microenvironment modeling. *Cell* **189**, 386–400 (2025).
78. Song, B. & Liang, R. Integrating artificial intelligence with smartphone-based imaging for cancer detection in vivo. *Biosens. Bioelectron.* **271**, 116982 (2025).
79. Shah, S. F. H. et al. Ethical implications of artificial intelligence in skin cancer diagnostics: use-case analyses. *Br. J. Dermatol.* **192**, 520–529 (2025).
80. Xu, H. et al. Artificial intelligence-assisted colonoscopy for colorectal cancer screening: a multicenter randomized controlled trial. *Clin. Gastroenterol. Hepatol.* **21**, 337–346.e3 (2023).
81. Ahmad, A. et al. Evaluation of a real-time computer-aided polyp detection system during screening colonoscopy: AI-DETECT study. *Endoscopy* **55**, 313–319 (2023).
82. Ardila, D. et al. End-to-end lung cancer screening with three-dimensional deep learning on low-dose chest computed tomography. *Nat. Med.* **25**, 954–961 (2019).
83. Christiansen, F. et al. International multicenter validation of AI-driven ultrasound detection of ovarian cancer. *Nat. Med.* **31**, 189–196 (2025).
84. Eisemann, N. et al. Nationwide real-world implementation of AI for cancer detection in population-based mammography screening. *Nat. Med.* **31**, 917–924 (2025).
85. Hamm, C. A. et al. Interactive explainable deep learning model informs prostate cancer diagnosis at MRI. *Radiology* **307**, e222276 (2023).
86. Haue, A. D., Hjaltelin, J. X., Holm, P. C., Placido, D. & Brunak, S. R. Artificial intelligence-aided data mining of medical records for cancer detection and screening. *Lancet Oncol.* **25**, e694–e703 (2024).
87. Hernström, V. et al. Screening performance and characteristics of breast cancer detected in the Mammography Screening with Artificial Intelligence trial (MASAI): a randomised, controlled, parallel-group, non-inferiority, single-blinded, screening accuracy study. *Lancet Digit. Health* **7**, e175–e183 (2025).
88. Jones, O. T. et al. Artificial intelligence and machine learning algorithms for early detection of skin cancer in community and primary care settings: a systematic review. *Lancet Digit. Health* **4**, e466–e476 (2022).
89. Lu, M. Y. et al. AI-based pathology predicts origins for cancers of unknown primary. *Nature* **594**, 106–110 (2021).
This work demonstrates a powerful clinical application of AI in pathology, showing that a deep learning model can accurately predict the tissue of origin for metastatic tumours, a key challenge in managing cancer of unknown primary.
90. Lång, K. et al. Artificial intelligence-supported screen reading versus standard double reading in the Mammography Screening with Artificial Intelligence trial (MASAI): a clinical safety analysis of a randomised, controlled, non-inferiority, single-blinded, screening accuracy study. *Lancet Oncol.* **24**, 936–944 (2023).
91. Karsenti, D. et al. Effect of real-time computer-aided detection of colorectal adenoma in routine colonoscopy (COLO-GENIUS): a single-centre randomised controlled trial. *Lancet Gastroenterol. Hepatol.* **8**, 726–734 (2023).
92. Ortiz, O. et al. An artificial intelligence-assisted system versus white light endoscopy alone for adenoma detection in individuals with Lynch syndrome (TIMELY): an international, multicentre, randomised controlled trial. *Lancet Gastroenterol. Hepatol.* **9**, 802–810 (2024).
93. Yuan, X. L. et al. Effect of an artificial intelligence-assisted system on endoscopic diagnosis of superficial oesophageal squamous cell carcinoma and precancerous lesions: a multicentre, tandem, double-blind, randomised controlled trial. *Lancet Gastroenterol. Hepatol.* **9**, 34–44 (2024).
94. Steyaert, S. et al. Multimodal data fusion for cancer biomarker discovery with deep learning. *Nat. Mach. Intell.* **5**, 351–362 (2023).
95. Zhu, G. et al. A deep-learning model for quantifying circulating tumour DNA from the density distribution of DNA-fragment lengths. *Nat. Biomed. Eng.* **9**, 307–319 (2025).
96. Cohen, J. D. et al. Detection and localization of surgically resectable cancers with a multi-analyte blood test. *Science* **359**, 926–930 (2018).
97. Liu, M. C., Oxnard, G. R., Klein, E. A., Swanton, C. & Seiden, M. V. Sensitive and specific multi-cancer detection and localization using methylation signatures in cell-free DNA. *Ann. Oncol.* **31**, 745–759 (2020).
98. Yang, M. et al. Enhancing the differential diagnosis of small pulmonary nodules: a comprehensive model integrating plasma methylation, protein biomarkers, and LDCT imaging features. *J. Transl. Med.* **22**, 984 (2024).
99. Guerra, C. E., Sharma, P. V. & Castillo, B. S. Multi-cancer early detection: the new frontier in cancer early detection. *Annu. Rev. Med.* **75**, 67–81 (2024).
100. Schrag, D. et al. Blood-based tests for multicancer early detection (PATHFINDER): a prospective cohort study. *Lancet* **402**, 1251–1260 (2023).
101. Bai, L. et al. Cancer biomarkers discovered using pan-cancer plasma proteomic profiling. *Nat. Biomed. Eng.* **10**, 16–38 (2026).
102. Malara, N. et al. Multicancer screening test based on the detection of circulating non haematological proliferating atypical cells. *Mol. Cancer* **23**, 32 (2024).
103. Zeune, L. L. et al. Deep learning of circulating tumour cells. *Nat. Mach. Intell.* **2**, 124–133 (2020).
104. Foroutan, F. et al. Computer aided detection and diagnosis of polyps in adult patients undergoing colonoscopy: a living clinical practice guideline. *BMJ* **388**, e082656 (2025).
105. Djinbachian, R. et al. Autonomous artificial intelligence vs artificial intelligence-assisted human optical diagnosis of colorectal polyps: a randomized controlled trial. *Gastroenterology* **167**, 392–399 e392 (2024).
106. Ebigbo, A., Messmann, H. & Lee, S. H. Artificial intelligence applications in image-based diagnosis of early esophageal and gastric neoplasms. *Gastroenterology* **169**, 396–415.e2 (2025).
107. Ding, W. et al. Contrast-enhanced ultrasound-based AI model for multi-classification of focal liver lesions. *J. Hepatol.* **83**, 426–439 (2025).
108. Lam, S. et al. Current and future perspectives on computed tomography screening for lung cancer: a roadmap from 2023 to 2027 from the International Association for the Study of Lung Cancer. *J. Thorac. Oncol.* **19**, 36–51 (2024).
109. Lam, S. et al. The International Association for the Study of Lung Cancer Early Lung Imaging Confederation Open-source Deep Learning and Quantitative Measurement Initiative. *J. Thorac. Oncol.* **19**, 94–105 (2024).
110. Darmofal, M. et al. Deep-Learning model for tumor-type prediction using targeted clinical genomic sequencing data. *Cancer Discov.* **14**, 1064–1081 (2024).
111. Chen, L. et al. Machine learning predicts oxaliplatin benefit in early colon cancer. *J. Clin. Oncol.* **42**, 1520–1530 (2024).
112. Adeoye, J. & Su, Y. X. Leveraging artificial intelligence for perioperative cancer risk assessment of oral potentially malignant disorders. *Int. J. Surg.* **110**, 1677–1686 (2024).

113. Chen, Z. et al. Predicting gastric cancer response to anti-HER2 therapy or anti-HER2 combined immunotherapy based on multi-modal data. *Signal. Transduct. Target. Ther.* **9**, 222 (2024).
114. Yang, Z. et al. A foundation model for generalizable cancer diagnosis and survival prediction from histopathological images. *Nat. Commun.* **16**, 2366 (2025).
115. Mikhael, P. G. et al. Sybil: a validated deep learning model to predict future lung cancer risk from a single low-dose chest computed tomography. *J. Clin. Oncol.* **41**, 2191–2200 (2023).
116. Yala, A. et al. Multi-institutional validation of a mammography-based breast cancer risk model. *J. Clin. Oncol.* **40**, 1732–1740 (2022).
117. Eriksson, M., Czene, K., Vachon, C., Conant, E. F. & Hall, P. Long-term performance of an image-based short-term risk model for breast cancer. *J. Clin. Oncol.* **41**, 2536–2545 (2023).
118. Huang, K. et al. A foundation model for clinician-centered drug repurposing. *Nat. Med.* **30**, 3601–3613 (2024).
119. Lapi, S. E. et al. Recent advances and impending challenges for the radiopharmaceutical sciences in oncology. *Lancet Oncol.* **25**, e236–e249 (2024).
120. Moon, I. et al. Machine learning for genetics-based classification and treatment response prediction in cancer of unknown primary. *Nat. Med.* **29**, 2057–2067 (2023).
121. Ogier du Terrail, J. et al. Federated learning for predicting histological response to neoadjuvant chemotherapy in triple-negative breast cancer. *Nat. Med.* **29**, 135–146 (2023).
This study provides a crucial proof-of-concept for federated learning, showing how multiple institutions can collaboratively train a powerful predictive model without sharing sensitive patient data.
122. Sammut, S.-J. et al. Multi-omic machine learning predictor of breast cancer therapy response. *Nature* **601**, 623–629 (2022).
123. Bao, X. et al. A multiomics analysis-assisted deep learning model identifies a macrophage-oriented module as a potential therapeutic target in colorectal cancer. *Cell Rep. Med.* **5**, 101399 (2024).
124. Loupy, A. et al. Reshaping transplantation with AI, emerging technologies and xenotransplantation. *Nat. Med.* **31**, 2161–2173 (2025).
125. Fraunhoffer, N. et al. Development and validation of AI-assisted transcriptomic signatures to personalize adjuvant chemotherapy in patients with pancreatic ductal adenocarcinoma. *Ann. Oncol.* **35**, 780–791 (2024).
126. Zhang, K. et al. Artificial intelligence in drug development. *Nat. Med.* **31**, 45–59 (2025).
127. Park, S. et al. A deep learning model of tumor cell architecture elucidates response and resistance to CDK4/6 inhibitors. *Nat. Cancer* **5**, 996–1009 (2024).
128. Keyl, J. et al. Decoding pan-cancer treatment outcomes using multimodal real-world data and explainable artificial intelligence. *Nat. Cancer* **6**, 307–322 (2025).
129. L'Imperio, V. et al. Pathologist validation of a machine learning-derived feature for colon cancer risk stratification. *JAMA Netw. Open* **6**, e2254891 (2023).
130. Heckenbach, I. et al. Deep learning assessment of senescence-associated nuclear morphologies in mammary tissue from healthy female donors to predict future risk of breast cancer: a retrospective cohort study. *Lancet Digit. Health* **6**, e681–e690 (2024).
131. Spratt, D. E. et al. Artificial intelligence predictive model for hormone therapy use in prostate cancer. *NEJM Evid.* **2**, EVIDoa2300023 (2023).
132. Esteva, A. et al. Prostate cancer therapy personalization via multi-modal deep learning on randomized phase III clinical trials. *NPJ Digit. Med.* **5**, 71 (2022).
133. Foersch, S. et al. Multistain deep learning for prediction of prognosis and therapy response in colorectal cancer. *Nat. Med.* **29**, 430–439 (2023).
134. Amgad, M. et al. A population-level digital histologic biomarker for enhanced prognosis of invasive breast cancer. *Nat. Med.* **30**, 85–97 (2024).
This study presents a population-level, deep learning-based biomarker from H&E slides that consistently outperforms pathologists in predicting breast cancer survival, highlighting the power of AI to capture prognostic stromal and immune features.
135. Wulczyn, E. et al. Interpretable survival prediction for colorectal cancer using deep learning. *NPJ Digit. Med.* **4**, 71 (2021).
136. Mataraso, S. J. et al. A machine learning approach to leveraging electronic health records for enhanced omics analysis. *Nat. Mach. Intell.* **7**, 293–306 (2025).
137. Volinsky-Fremond, S. et al. Prediction of recurrence risk in endometrial cancer with multimodal deep learning. *Nat. Med.* **30**, 1962–1973 (2024).
This study demonstrates that a multi-modal deep learning model (HECTOR), integrating histology and clinical stage, outperforms the current molecular-based gold standard for predicting distant recurrence in endometrial cancer.
138. Kehl, K. L. et al. Shareable artificial intelligence to extract cancer outcomes from electronic health records for precision oncology research. *Nat. Commun.* **15**, 9787 (2024).
139. Gale, D. et al. Residual ctDNA after treatment predicts early relapse in patients with early-stage non-small cell lung cancer. *Ann. Oncol.* **33**, 500–510 (2022).
140. Pantel, K. & Alix-Panabières, C. Minimal residual disease as a target for liquid biopsy in patients with solid tumours. *Nat. Rev. Clin. Oncol.* **22**, 65–77 (2025).
141. Bruhm, D. C. et al. Genomic and fragmentomic landscapes of cell-free DNA for early cancer detection. *Nat. Rev. Cancer* **25**, 341–358 (2025).
142. Abbosh, C. et al. Tracking early lung cancer metastatic dissemination in TRACERx using ctDNA. *Nature* **616**, 553–562 (2023).
143. Wang, Y. et al. PRIME: an interpretable artificial intelligence model based on liquid biopsy improves prediction of progression risk in non-small cell lung cancer. *Mil. Med. Res.* **12**, 94 (2026).
144. Loeffler, C. M. L. et al. HIBRID: histology-based risk-stratification with deep learning and ctDNA in colorectal cancer. *Nat. Commun.* **16**, 7561 (2025).
145. Thangaraj, P. M., Benson, S. H., Oikonomou, E. K., Asselbergs, F. W. & Khera, R. Cardiovascular care with digital twin technology in the era of generative artificial intelligence. *Eur. Heart J.* **45**, 4808–4821 (2024).
146. Zhang, K. et al. Concepts and applications of digital twins in healthcare and medicine. *Patterns* **5**, 101028 (2024).
147. Asghar, U. S. & Chung, C. Application of digital twins for personalized oncology. *Nat. Rev. Cancer* **25**, 823–825 (2025).
148. Sadée, C. et al. Medical digital twins: enabling precision medicine and medical artificial intelligence. *Lancet Digit. Health* **7**, 100864 (2025).
149. Reicher, L. et al. Deep phenotyping of health-disease continuum in the Human Phenotype Project. *Nat. Med.* **31**, 3191–3203 (2025).
150. National Academies of Sciences, Engineering & Medicine. *Foundational Research Gaps and Future Directions for Digital Twins* (The National Academies Press, 2024).
151. McCoy, M. Digital twins in oncology: where we are and where we hope to go. *BMJ Oncol.* **4**, e000893 (2025).
152. Gilbertson, R. J. et al. The virtual child. *Cancer Discov.* **14**, 663–668 (2024).
153. Laubenbacher, R., Mehrad, B., Shmulevich, I. & Trayanova, N. Digital twins in medicine. *Nat. Comput. Sci.* **4**, 184–191 (2024).
154. Asciak, L. et al. Digital twin assisted surgery, concept, opportunities, and challenges. *NPJ Digit. Med.* **8**, 32 (2025).
155. Ector, C. et al. Time-of-day effects of cancer drugs revealed by high-throughput deep phenotyping. *Nat. Commun.* **15**, 7205 (2024).
156. Foy, B. H. et al. Haematological setpoints are a stable and patient-specific deep phenotype. *Nature* **637**, 430–438 (2025).
157. Görtz, M. et al. Digital twins for personalized treatment in uro-oncology in the era of artificial intelligence. *Nat. Rev. Urol.* **23**, 29–39 (2025).
158. Tushar, F. I. et al. Virtual lung screening trial (VLST): an in silico study inspired by the national lung screening trial for lung cancer detection. *Med. Image Anal.* **103**, 103576 (2025).
159. Sarrami-Foroushani, A. et al. In-silico trial of intracranial flow diverters replicates and expands insights from conventional clinical trials. *Nat. Commun.* **12**, 3861 (2021).
160. Scibilia, K. R. et al. Mathematical oncology: how modeling is transforming clinical decision-making. *Cancer Res.* **85**, 4866–4879 (2025).
161. Demir, S. et al. Mechanistic models position ceritinib as a nuclear integrity disrupting therapy in pediatric liver tumors. *J. Exp. Clin. Cancer Res.* **44**, 268 (2025).
162. DeGrave, A. J., Cai, Z. R., Janizek, J. D., Daneshjou, R. & Lee, S. I. Auditing the inference processes of medical-image classifiers by leveraging generative AI and the expertise of physicians. *Nat. Biomed. Eng.* **9**, 294–306 (2025).
163. Rong, Y. et al. Towards human-centered explainable AI: a survey of user studies for model explanations. *IEEE Trans. Pattern Anal. Mach. Intell.* **46**, 2104–2122 (2024).
164. Chen, D. et al. Physician and artificial intelligence chatbot responses to cancer questions from social media. *JAMA Oncol.* **10**, 956–960 (2024).
165. Benfatto, S. et al. Explainable artificial intelligence of DNA methylation-based brain tumor diagnostics. *Nat. Commun.* **16**, 1787 (2025).
This work pioneers the concept of a 'biologically informed' neural network, which integrates known biological pathways into the model architecture to improve prediction and enable the discovery of novel molecular drivers of treatment resistance.
166. Ghassemi, M., Oakden-Rayner, L. & Beam, A. L. The false hope of current approaches to explainable artificial intelligence in health care. *Lancet Digit. Health* **3**, e745–e750 (2021).
167. Feuerriegel, S. et al. Causal machine learning for predicting treatment outcomes. *Nat. Med.* **30**, 958–968 (2024).
168. Elmarakeby, H. A. et al. Biologically informed deep neural network for prostate cancer discovery. *Nature* **598**, 348–352 (2021).
169. Novakovsky, G., Dexter, N., Libbrecht, M. W., Wasserman, W. W. & Mostafavi, S. Obtaining genetics insights from deep learning via explainable artificial intelligence. *Nat. Rev. Genet.* **24**, 125–137 (2023).
170. Chen, H., Gomez, C., Huang, C. M. & Unberath, M. Explainable medical imaging AI needs human-centered design: guidelines and evidence from a systematic review. *NPJ Digit. Med.* **5**, 156 (2022).
171. Kurian, M., Adashek, J. J. & West, H. J. Cancer care in the era of artificial intelligence. *JAMA Oncol.* **10**, 683 (2024).
172. Acosta, J. N., Falcone, G. J., Rajpurkar, P. & Topol, E. J. Multimodal biomedical AI. *Nat. Med.* **28**, 1773–1784 (2022).
173. Alderman, J. E. et al. Tackling algorithmic bias and promoting transparency in health datasets: the STANDING together consensus recommendations. *Lancet Digit. Health* **7**, e64–e88 (2025).
174. Fiske, A. et al. Weighing the benefits and risks of collecting race and ethnicity data in clinical settings for medical artificial intelligence. *Lancet Digit. Health* **7**, e286–e294 (2025).
175. Giddings, R. et al. Factors influencing clinician and patient interaction with machine learning-based risk prediction models: a systematic review. *Lancet Digit. Health* **6**, e131–e144 (2024).
176. Gazzarata, R. et al. HL7 Fast Healthcare Interoperability Resources (HL7 FHIR) in digital healthcare ecosystems for chronic disease management: scoping review. *Int. J. Med. Inf.* **189**, 105507 (2024).
177. Vokinger, K. N. & Gasser, U. Regulating AI in medicine in the United States and Europe. *Nat. Mach. Intell.* **3**, 738–739 (2021).

178. Qian, X. et al. A multimodal machine learning model for the stratification of breast cancer risk. *Nat. Biomed. Eng.* **9**, 356–370 (2025).
179. Kehayias, C. E. et al. A prospectively deployed deep learning-enabled automated quality assurance tool for oncological palliative spine radiation therapy. *Lancet Digit. Health* **7**, e13–e22 (2025).
180. Moons, K. G. M. et al. PROBAST+AI: an updated quality, risk of bias, and applicability assessment tool for prediction models using regression or artificial intelligence methods. *BMJ* **388**, e082505 (2025).

Acknowledgements
This research was supported by the National Natural Science Foundation of China (32470964 and 32100631 to F.L., W2431057 to K.Z.), the Macau Science and Technology Development Fund, Macao (0007/2020/AFJ, 0070/2020/A2, and 0003/2021/AKP to K.Z.), and the Guangzhou National Laboratory (YW-SLJC0201 to K.Z.), sponsored by Beijing Nova Program (20240484627 to F.L.), Capital's Funds for Health Improvement and Research (CFH2024-4-40215 to F.L.), China Postdoctoral Science Foundation (2023T160061 to F.L.), and Macao Young Scholars Program (AM2023018 to F.L.).

Author contributions
All authors researched data for the article, contributed substantially to discussion of the content, and reviewed and/or edited the manuscript before submission. F.L., H.L. and K.Z. wrote the article.

Competing interests
The authors declare no competing interests.

Additional information
Peer review information *Nature Reviews Cancer* thanks the anonymous reviewers for their contribution to the peer review of this work.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.