

Review Article

Can Liquid Biopsy Predict the Future of a Tumour? A Critical Review of AI-Enabled Longitudinal Modelling for Cancer Evolution and Treatment Resistance

Richard Afriyie Osei¹, Danmegoro Suleman Taiye², Godwin Kenechukwu Nwajiugo³, Oluwabukunmi Oluwaniran⁴, Ismaila Muhammed⁵

¹Department of Nanoscience, University of North Carolina at Greensboro, Greensboro, NC, USA

²ER Department, Ar-Rafiah General Hospital, Ministry of Health, Dammam, Saudi Arabia

³Department of Radiotherapy and Clinical Oncology, National Hospital Abuja, Abuja, Nigeria

⁴Department of Anatomic Pathology, Federal Medical Center, Abeokuta, Nigeria

⁵Department of Mathematics, Khalifa University, Abu Dhabi, United Arab Emirates

Received: Sep 07, 2025

Accepted: Sep 27, 2026

Corresponding author's email:

richoy22@gmail.com

Citation: Osei RA, Taiye DS, Nwajiugo GK, Oluwaniran O, Muhammed I. Can Liquid Biopsy Predict the Future of a Tumour? A Critical Review of AI-Enabled Longitudinal Modelling for Cancer Evolution and Treatment Resistance. *Oncology, Nuclear Medicine and Transplantology*. 2026;2(4):onmt029.

<https://doi.org/10.63946/onmt/19502>

Copyright: By the author(s).

License: A non-exclusive license by the publisher. Published by Australasia Publishing Group LLP (Astana, Kazakhstan) on behalf National Research Oncology Center.

Open Access: This article is an open access article distributed under the terms and conditions of the CC-BY Creative Commons Attribution license

<https://creativecommons.org/licenses/by/4.0>



ABSTRACT

Cancer behaves as a dynamic evolutionary system, whereas conventional tissue biopsy and imaging provide incomplete temporal information about how a tumour is changing. This review critically examines whether longitudinal liquid-biopsy measurement, combined with artificial intelligence (AI), can move oncology from tumour monitoring toward genuine prediction of cancer evolution and treatment resistance. Using a narrative critical-review approach, representative mechanistic, clinical, and computational evidence was purposively synthesised to evaluate the current maturity of AI-enabled longitudinal liquid-biopsy prediction. We synthesise evidence on circulating tumour DNA, circulating tumour cells, extracellular vesicles, and multimodal liquid-biopsy platforms, together with AI approaches applied to longitudinal cancer data, drawing on landmark and recent studies across colorectal, lung, breast, and ovarian cancer. Current evidence distinguishes three increasingly demanding tasks: early detection, longitudinal monitoring, and genuine forecasting. Longitudinal circulating tumour DNA analysis can detect molecular progression and emerging resistance mutations months before radiographic progression, and machine-learning models can stratify survival using serial biomarker metrics; however, most published “AI prediction” studies demonstrate earlier detection or prognostic association rather than prospectively validated forecasting of a future, patient-specific event. Biological variability in tumour shedding, analytical noise, small and irregularly sampled longitudinal datasets, overfitting, data leakage, weak external validation, and limited interpretability constrain the transition from detection to actionable prediction. AI-enabled liquid biopsy has genuine potential to forecast tumour trajectories, but current evidence is considerably stronger for molecular monitoring and early detection than for prospectively validated, clinically actionable prediction. Realising this potential will require standardised assays, prospective longitudinal cohorts with predefined prediction horizons, rigorous external and temporal validation, calibrated and interpretable models, and interventional trials linking predictions to treatment decisions before AI-enabled liquid biopsy can be considered ready to guide clinical care.

Keywords: Liquid Biopsy; Circulating Tumour DNA; Artificial Intelligence; Machine Learning; Tumour Evolution; Treatment Resistance; Longitudinal Monitoring; Precision Oncology

Introduction

The problem: cancer is an evolving system

Cancer is not a static entity but a dynamic evolutionary system in which genetically and phenotypically diverse cellular populations continually emerge, compete, and adapt [1]. Genomic instability generates intratumoural heterogeneity, providing a substrate for clonal selection and disease progression [2]. Therapeutic intervention further reshapes this ecosystem: treatment can eliminate susceptible populations while imposing selective pressure that favours pre-existing or newly acquired resistant subclones [3]. Consequently, the molecular composition of a tumour may change substantially during treatment, making a single baseline characterisation increasingly inadequate for understanding subsequent disease behaviour.

Tissue biopsy is particularly limited in this context because it samples a restricted anatomical site at a single time point and therefore may not capture spatially distributed or temporally emerging clones. Serial plasma analysis offers a complementary strategy by enabling repeated, minimally invasive assessment of tumour-derived molecular signals and has demonstrated the capacity to track clonal evolution and treatment-associated resistance [4,5].

Why liquid biopsy changes the problem

Liquid biopsy introduces a fundamentally different sampling paradigm by enabling repeated, minimally invasive interrogation of tumour-derived material in circulation. Circulating tumour DNA (ctDNA) provides genomic information on tumour burden, clonal composition, and emerging molecular alterations [6], while circulating tumour cells (CTCs) can retain cellular and phenotypic features that complement DNA-based measurements [7]. Extracellular vesicles and circulating RNA provide additional information on intercellular signalling and gene regulation [10], whereas DNA methylation and fragmentomic patterns can capture tissue-of-origin and tumour-associated epigenetic features. Circulating proteins may further contribute functional and phenotypic information. Rather than treating these analytes as interchangeable biomarkers, their complementary biological signals create an opportunity to construct multidimensional molecular trajectories. Multi-analyte platforms that integrate several circulating signals simultaneously have already improved diagnostic performance relative to single-analyte tests [12], and recent advances increasingly emphasise integrated, multi-omic liquid-biopsy profiling, although analytical variability, low tumour-derived signal, and incomplete biological representation remain important constraints.

Why AI is entering the field

The increasing use of liquid biopsy in longitudinal cancer monitoring has created a computational problem as much as a biological one. Serial samples can contain numerous molecular features measured at different time points, with signals that vary between patients and fluctuate during disease progression and treatment. These data are therefore high-dimensional, heterogeneous, and inherently time-dependent, while low concentrations of tumour-derived material and technical variation can introduce substantial noise. Interpreting such patterns manually becomes increasingly difficult as the number of measurements and biological variables expands. Machine learning and deep learning provide a means of analysing these complex relationships simultaneously, identifying patterns across repeated measurements and potentially distinguishing clinically meaningful trajectories from random variation [13]. Recent studies demonstrate the growing application of machine learning to liquid-biopsy data, particularly where multiple molecular features must be integrated to improve prediction [17].

The critical question

The central question is no longer whether liquid biopsy can provide information about cancer, but whether changes in circulating biomarkers can be interpreted reliably enough to anticipate what the tumour will do next. This distinction is important because an earlier molecular signal of progression or resistance does not, by itself, constitute a prediction of future tumour behaviour. AI may provide the analytical capacity required to model complex trajectories across serial measurements, but the clinical value of such models depends on whether they can generalise beyond the datasets in which they were developed and provide sufficiently early, accurate, and actionable forecasts. Current evidence increasingly supports longitudinal monitoring and early identification of molecular changes, while prospective evidence demonstrating that AI models can reliably forecast tumour evolution and guide treatment before clinical progression remains comparatively limited.

This review therefore critically evaluates whether AI-enabled longitudinal analysis of liquid-biopsy biomarkers can move cancer monitoring beyond retrospective detection and association towards prospective prediction of tumour evolution and treatment resistance. Particular attention is given to the distinction between monitoring, early detection, and genuine forecasting, the biological and computational factors that constrain prediction, and the level of

validation required before such models can inform clinical decisions.

Review Approach

This article was conducted as a narrative critical review rather than a systematic review. Accordingly, no formal systematic database search, predefined search strategy, or systematic study-selection protocol was undertaken. The literature was purposively identified and selected to provide representative mechanistic, clinical, and computational evidence relevant to longitudinal liquid biopsy, tumour

evolution, treatment resistance, and AI-enabled prediction. Particular emphasis was placed on landmark studies and recent evidence illustrating serial liquid-biopsy monitoring, emerging resistance, longitudinal modelling, and the methodological distinction between early molecular detection, prognostic association, and genuine prospective forecasting. The selected literature was critically synthesised to evaluate the current maturity, limitations, and clinical translation of AI-enabled longitudinal liquid-biopsy approaches.

From Static Biomarkers to Dynamic Tumour Monitoring

The limitations of conventional monitoring

Conventional cancer monitoring provides valuable information, but much of it remains inherently limited in space or time. A tissue biopsy offers detailed histological and molecular information, yet it represents a sample from a particular anatomical site at a particular moment. In heterogeneous tumours, this spatial snapshot may not reflect genetically distinct populations elsewhere in the primary tumour or at metastatic sites, while repeated tissue sampling is often impractical because of invasiveness and anatomical constraints.

Imaging provides a different perspective by tracking anatomical changes during treatment, but changes in tumour size or morphology do not necessarily correspond directly to the underlying molecular events driving progression or resistance. A single liquid biopsy similarly provides only a molecular snapshot. Its greater potential emerges when blood samples are collected repeatedly, allowing changes in circulating tumour-derived signals to be examined as a trajectory rather than as isolated measurements. Longitudinal sampling can therefore capture temporal changes that a single tissue specimen or imaging assessment may miss, creating an opportunity to observe tumour evolution as it unfolds during treatment.

Longitudinal liquid biopsy as a window into tumour evolution

Serial liquid biopsy provides an opportunity to observe molecular changes as treatment reshapes the tumour. Changes in ctDNA abundance can reflect alterations in tumour burden, while the appearance, disappearance, or changing allele frequency of specific variants may indicate shifts in the relative abundance of tumour clones. Over successive sampling points, these patterns can reveal the expansion of resistant populations that may be difficult to identify from a single biopsy. In metastatic colorectal cancer, longitudinal ctDNA progression preceded

radiographic progression in 58.1% of patients, with a median lead time of 3.3 months, and resistant mutations and gene amplifications were identified at the time of progression [16].

The value of longitudinal analysis, however, extends beyond detecting individual mutations. The direction and rate of molecular change may provide information about treatment response and the trajectory of disease. Studies in metastatic lung and breast cancer have similarly shown associations between serial ctDNA dynamics, treatment response, and subsequent clinical outcomes [17,20]. Persistent postoperative ctDNA is likewise associated with a substantially increased risk of recurrence, providing a lead time over standard surveillance imaging in several solid tumours [22,23]. These findings support longitudinal liquid biopsy as a means of observing tumour evolution between conventional clinical assessments.

But detection is not prediction

For clarity, four related but distinct terms are used throughout this review. Detection refers to identifying the presence of a tumour-associated molecular signal at a given time point. Monitoring refers to serial measurement of that signal to determine how it changes over time. Prognosis refers to the estimation of a patient's subsequent clinical outcome or risk based on information available at a defined assessment point; although prognostic associations may identify patients at higher risk of progression, recurrence, or death, they do not necessarily constitute longitudinal forecasting of how a tumour will evolve over time. Forecasting refers to using information available at an earlier, predefined prediction point to estimate the probability or timing of a future event that has not yet occurred. Actionable clinical prediction represents a higher threshold in which a forecast has sufficient accuracy, validation, lead time, and clinical relevance to support a treatment decision. These terms are not interchangeable: an early molecular signal that precedes radiographic or clinical progression

represents early detection or monitoring unless it is prospectively used to estimate a subsequent event within a defined prediction horizon.

This distinction is particularly important in longitudinal liquid-biopsy research. A rise in ctDNA that precedes radiographic progression may demonstrate earlier detection of disease activity, but it does not necessarily establish that a model can predict future progression in an individual patient. Similarly, identifying a resistance-associated mutation before treatment failure demonstrates molecular evolution but

does not by itself establish when resistance will become clinically consequential, or whether an alternative treatment would improve the outcome. Longitudinal liquid biopsy should therefore be viewed as a continuum from detection, to monitoring, to forecasting, and ultimately to actionable prediction. The critical challenge for AI is to demonstrate that it can make the transition between these stages reliably, rather than simply classify patterns that are already apparent.

What Liquid Biopsy Can Actually Tell Us About Tumour Evolution

ctDNA and clonal evolution

Among liquid-biopsy biomarkers, ctDNA provides one of the most direct means of following the genomic dynamics of a tumour over time. Its value is not simply that tumour-associated variants can be detected in plasma [6], but that their abundance can be measured repeatedly. Changes in variant allele frequency across serial samples may therefore provide an indirect view of how individual clones expand or contract during disease progression and treatment. A declining allele frequency may accompany the suppression of a treatment-sensitive clone, whereas persistence or increasing representation of another variant may indicate selection of a population that has become less responsive to therapy.

This longitudinal perspective is particularly relevant to treatment resistance [32]. Resistant subclones may exist before therapy and subsequently become more prominent under treatment pressure. Recent single-cell and cell-free DNA sequencing in high-grade serous ovarian cancer showed that drug-resistant disease typically arises from the selective expansion of a small subset of clones already present at diagnosis, and that this expansion could be tracked non-invasively in serial plasma samples across the treatment course [21]. Similar applications of ctDNA include identifying molecular relapse and detecting minimal residual disease after apparently successful treatment. Persistent ctDNA after definitive therapy is associated with a substantially increased risk of recurrence in several solid tumours, most robustly demonstrated in resected stage II colon cancer, although the strength of this association varies with tumour type, assay, and clinical setting [22,23].

These observations establish ctDNA as a useful marker of tumour dynamics, but they also expose an important limitation. A changing molecular signal indicates that the tumour population is changing; it does not necessarily reveal the complete evolutionary trajectory or establish which future clinical event will

occur. Low tumour shedding, biological heterogeneity, and technical differences between assays can also obscure genuine changes or produce measurements that are difficult to compare across time. Thus, the principal opportunity for AI is not merely to identify individual mutations, but to determine whether complex patterns of clonal change across repeated measurements contain sufficiently reliable information to forecast subsequent tumour behaviour.

CTCs and cellular heterogeneity

CTCs provide information that cannot be fully recovered from cell-free DNA because they are intact tumour cells rather than molecular fragments released into the circulation [7]. Their analysis can therefore preserve aspects of cellular phenotype, including protein expression, morphology, and transcriptional state, while also allowing direct investigation of genomic features at the level of individual cells. This distinction is important in heterogeneous tumours, where clinically relevant subpopulations may differ substantially in their capacity for proliferation, invasion, survival, and treatment response.

CTCs may also provide a more direct view of metastatic biology. Their presence in the circulation represents a stage in the dissemination process, and specific CTC phenotypes and clusters have been associated with enhanced metastatic capacity [8]. Recent phylogenetic analysis using whole-exome sequencing has further shown that individual CTC clusters can be genetically oligoclonal, containing cells derived from more than one tumour clone, highlighting their potential role in polyclonal dissemination and tumour evolution [9].

However, the biological richness of CTCs is accompanied by a major analytical limitation: they are extremely rare, and their isolation depends strongly on the characteristics of the enrichment and detection method. Consequently, CTC measurements may be less reproducible and less readily scalable than ctDNA analysis. Their greatest value may therefore lie not in replacing ctDNA, but in complementing it with cellular

and phenotypic information that can help explain the biological significance of molecular changes observed in plasma.

Extracellular vesicles and circulating RNA

Extracellular vesicles and circulating RNA provide a complementary view of tumour biology because they can reflect active communication between malignant cells and their surrounding environment. Extracellular vesicles carry biologically active cargo, including RNA, DNA, proteins, and lipids, and can transfer this material between tumour cells and stromal or immune cells [10]. This communication can alter signalling pathways involved in tumour growth, immune regulation, and adaptation to treatment, thereby contributing to the development of therapeutic resistance.

Circulating RNA, particularly RNA protected within extracellular vesicles, may capture changes in gene regulation that are not readily apparent from genomic analysis alone. Extracellular-vesicle-derived long RNA signatures have, for example, been combined with computed tomography (CT) features to improve discrimination of early-stage lung adenocarcinoma beyond either modality alone [11], supporting the potential value of vesicle-derived RNA for monitoring tumour biology and treatment response. However, the biological heterogeneity of extracellular vesicles and differences in their isolation and characterisation remain important limitations. These issues are particularly relevant when measurements are intended for longitudinal modelling, because a signal must be reproducible across time before changes can be confidently interpreted as biological rather than technical. Extracellular vesicle and circulating RNA data may therefore add mechanistic information to longitudinal liquid biopsy, but their predictive value will depend on analytical standardisation and validation in appropriately designed patient cohorts. Despite their biological promise, the evidence supporting CTCs, EVs, and circulating RNA for longitudinal forecasting remains substantially less mature than that for ctDNA. Much of the available literature has focused on analyte detection, molecular characterisation, diagnostic discrimination, or associations with prognosis rather than prospectively validated prediction of future tumour evolution or treatment resistance. Longitudinal studies with predefined prediction points, clinically relevant prediction horizons, and independent validation remain limited. These analytes should therefore be regarded as promising complementary sources of temporal information rather than established substrates for AI-enabled clinical forecasting.

Multi-analyte liquid biopsy

The biological complexity of tumour evolution raises a limitation of relying on any single circulating biomarker. ctDNA can provide genomic evidence of changing tumour populations, but it does not directly capture cellular phenotype. CTCs can provide cellular and single-cell information, yet their scarcity can restrict sampling depth. Extracellular vesicles and circulating RNA may reflect signalling and regulatory processes that are not represented fully by genomic measurements. Clinical variables and imaging provide additional information about the patient's disease state and its anatomical consequences. The challenge is therefore to determine whether these complementary signals can be combined into a coherent representation of tumour behaviour rather than interpreted independently.

Early studies provide evidence that multimodal approaches can improve biological classification and prognostic modelling. Combining liquid-biopsy measurements with CT features has improved diagnostic discrimination compared with individual modalities in early-stage lung adenocarcinoma [11], and multi-analyte blood tests that integrate protein and genomic markers have improved detection of surgically resectable cancers relative to single-marker approaches [12]. More recent work has similarly emphasised the potential of integrating circulating DNA, tumour cells, extracellular vesicles, proteins, and immune components to characterise tumour and host dynamics over time.

For AI-enabled longitudinal modelling, however, the relevant question is not simply whether more variables improve model performance. Each additional analyte introduces its own measurement error, biological variability, and missing data. A useful model must therefore determine which signals contain reproducible temporal information, which provide mechanistic context, and which contribute little beyond existing measurements. Multimodal integration is valuable only when the combined data improve prediction in independently validated patients and remain interpretable enough to support clinical decisions. The future challenge is consequently not to collect every available biomarker, but to identify the smallest and most reliable combination of longitudinal signals capable of representing tumour evolution with clinically meaningful accuracy. Although multimodal integration has demonstrated encouraging diagnostic and prognostic performance, such performance should not be equated with longitudinal forecasting. A model that combines molecular, imaging, and clinical variables to classify current disease status or estimate prognosis at a single time point does not necessarily predict how an individual tumour will evolve over

time. Evidence for multimodal models that integrate serial liquid-biopsy measurements with other longitudinal data to prospectively forecast treatment resistance, progression, or clonal evolution remains

limited. Future studies should therefore distinguish clearly between multimodal diagnostic or prognostic modelling and true longitudinal forecasting based on temporally ordered patient data.

AI-Enabled Longitudinal Modelling

From conventional longitudinal modelling to AI

Longitudinal liquid-biopsy data can be analysed using established statistical approaches, including mixed-effects models, longitudinal regression, survival analysis, and joint models that relate repeated biomarker measurements to time-to-event outcomes. These methods remain valuable because they provide interpretable estimates of temporal associations and can explicitly account for within-patient correlation, repeated measurements, and clinically relevant covariates. AI and machine-learning approaches may offer additional flexibility when datasets become high-dimensional, multimodal, or characterised by complex nonlinear relationships and interactions that are difficult to specify a priori. Time-series and deep-learning methods can further model dependencies across sequential observations and integrate molecular, imaging, treatment, and clinical features. However, greater algorithmic complexity does not inherently confer superior predictive performance or clinical utility. The relevant question is therefore not whether AI should replace established longitudinal statistical methods, but whether it provides reproducible incremental predictive value when evaluated against appropriate statistical benchmarks and validated in independent patient populations.

AI approaches relevant to tumour evolution

Different AI approaches address distinct aspects of modelling tumour evolution. **Supervised machine learning** is most directly applicable when a predefined clinical endpoint is available. Models can integrate liquid-biopsy features to estimate outcomes such as treatment response, progression, recurrence, or survival. For example, a longitudinal ctDNA model developed from a randomised phase 3 trial of 466 patients with non-small-cell lung cancer integrated multiple ctDNA metrics collected across five treatment time points using machine learning and stratified patients into groups with markedly different survival outcomes [17]. This study illustrates how machine learning can combine dynamic molecular features across serial measurements rather than relying on a single ctDNA value. However, survival-risk stratification represents prognostic modelling and is not equivalent to forecasting the specific biological

event or resistance mechanism responsible for subsequent progression. The intended prediction target and prediction horizon must therefore be considered when interpreting AI model performance.

Importantly, the performance evidence supporting AI-enabled longitudinal prediction remains heterogeneous. In the NSCLC study, the longitudinal ctDNA model identified patient groups with substantially different survival outcomes and retained prognostic value in an external cohort [17]. However, many other longitudinal liquid-biopsy studies report measures of molecular lead time, hazard ratios, or biomarker discrimination rather than the full set of metrics expected for clinical prediction models, including calibration, sensitivity, specificity, predictive values, and decision-curve analysis. Consequently, direct comparison of predictive performance across studies remains difficult, and evidence of association or discrimination should not be interpreted as demonstrating a well-calibrated and clinically useful forecasting model.

Unsupervised learning addresses a different problem because it does not require a predefined clinical outcome. Clustering and related approaches can identify molecularly distinct patient groups or patterns of longitudinal biomarker change that may correspond to different evolutionary states. Such methods may be particularly useful for discovering previously unrecognised trajectory patterns or molecular subgroups that could subsequently inform predictive modelling. However, these data-driven clusters do not themselves constitute predictions, and their biological and clinical relevance requires independent validation before they can be interpreted as meaningful evolutionary states.

Time-series and dynamic survival models are particularly relevant because tumour evolution is inherently temporal. Recurrent neural networks and related architectures, including long short-term memory networks, are designed to learn dependencies across sequential observations [14] and could therefore model evolving biomarker trajectories in relation to subsequent progression, treatment response, or resistance. Dynamic survival models address a related task by updating estimates of the risk or timing of events such as progression, relapse, or death as new longitudinal measurements become available. Transformer-based architectures use self-attention to

model relationships across sequences of observations [15] and may provide another approach for integrating multiple time-dependent molecular and clinical features. However, these architectures should be regarded as methodologically relevant approaches rather than established liquid-biopsy forecasting tools unless they have been specifically evaluated and validated using longitudinal liquid-biopsy data.

Multimodal AI extends longitudinal modelling by integrating liquid-biopsy measurements with complementary information such as imaging, treatment history, and clinical characteristics. Each modality captures a different component of disease: circulating biomarkers provide molecular information, imaging captures anatomical change, and clinical data provide patient and treatment context. Multimodal models may therefore be designed to predict outcomes such as diagnosis, treatment response, progression, or prognosis by combining signals that would otherwise be analysed separately [11]. Nevertheless, much of the existing multimodal evidence remains diagnostic or prognostic rather than demonstrating prospective longitudinal forecasting. Additional modalities also introduce missingness, measurement heterogeneity, and opportunities for overfitting. The critical test is therefore whether multimodal models provide reproducible improvements in temporally separated, prospective, or externally validated cohorts rather than simply achieving higher performance within their development datasets.

What would constitute genuine prediction?

For AI-enabled longitudinal modelling to constitute genuine prediction, information used by the model must be available before the outcome of interest and should be used to estimate a subsequent event, such as progression or treatment resistance, within a predefined prediction horizon. This temporal separation is essential to distinguish forecasting from retrospective classification. Credible forecasting also requires evaluation beyond the development dataset. Independent cohorts, temporal validation, and assessment of calibration are needed to establish whether predictions generalise to new patients [27–29]. Ultimately, the strongest evidence would come from prospective evaluation in which predictions are generated before clinical progression and linked to predefined management strategies. Accordingly, longitudinal liquid-biopsy studies should clearly specify the prediction point, prediction horizon, outcome, and intended clinical use. These methodological requirements provide the basis for evaluating whether longitudinal liquid-biopsy studies have progressed from molecular monitoring and prognostic association to genuine forecasting. The following section therefore examines the available clinical evidence, with particular attention to what serial liquid-biopsy measurements have demonstrated, whether predictive models have been prospectively validated, and how close current approaches are to supporting clinically actionable decisions.

Can We Really Predict Treatment Resistance?

Evidence that resistance leaves molecular footprints

Treatment resistance can leave detectable molecular traces in the circulation as tumour populations are reshaped by therapeutic selection. Under treatment pressure, sensitive clones may decline while pre-existing or newly acquired resistant populations gain a relative advantage. Serial ctDNA analysis can capture this process through changes in variant allele frequencies, the emergence of previously undetected alterations, and the increasing representation of resistant clones. Such changes have been observed across several tumour types and can provide molecular evidence of evolutionary selection before resistance becomes apparent on conventional imaging.

The dynamics of ctDNA can also reflect treatment response and subsequent relapse. In a longitudinal study of 62 patients with metastatic colorectal cancer receiving first-line chemotherapy, ctDNA clearance was independently associated with

longer progression-free survival, whereas ctDNA progression preceded radiographic progression in 58.1% of patients, with a median lead time of 3.3 months; diverse resistant mutations and gene amplifications were identified at the time of progression [16]. Similar longitudinal approaches have been used to detect molecular residual disease and identify patients at increased risk of recurrence after apparently successful treatment [22,23].

These findings support the biological plausibility of using circulating biomarkers to follow treatment-induced evolutionary change. They also reveal an important limitation: not all resistance is explained by detectable genetic alterations. Cancer cells can acquire drug-tolerant states through epigenetic, transcriptional, and phenotypic adaptation without a readily identifiable resistance mutation. Consequently, the absence of a resistance-associated alteration cannot be interpreted as evidence that resistance is absent. The molecular footprint captured in blood is therefore an incomplete representation of a changing tumour ecosystem, reinforcing the need for models that

integrate multiple longitudinal signals rather than relying on isolated genetic events.

Evidence supporting early prediction

Evidence from several tumour types suggests that longitudinal liquid biopsy can provide molecular information before conventional evidence of treatment failure. In EGFR-mutant non-small cell lung cancer, serial plasma monitoring by droplet digital PCR detected the resistance mutation EGFR T790M in 47% of patients who progressed on first- or second-generation tyrosine kinase inhibitors, with almost half of positive cases identified a median of 2.2 months before clinically apparent progression [18]. Building on this observation, the prospective randomised APPLE trial demonstrated that serial plasma T790M monitoring is feasible for guiding the sequencing of gefitinib and osimertinib, identifying molecular progression before RECIST-defined progression in 17% of patients and enabling an earlier switch to osimertinib with a favourable 18-month progression-free survival [19]. This supports the value of repeated molecular assessment as an earlier indicator of changing treatment sensitivity, although the timing and clinical significance of individual resistance alterations vary, and their detection does not necessarily establish when clinically evident resistance will develop.

Colorectal cancer provides particularly clear evidence for the temporal value of ctDNA. In the longitudinal cohort described above, ctDNA progression preceded radiographic progression in 58.1% of patients, with a median lead time of 3.3 months [16]. These findings demonstrate that molecular progression can become apparent before anatomical progression, although the study primarily established earlier detection rather than prospective prediction of a future event.

Breast cancer studies provide a similar but more nuanced picture. In a prospective cohort of patients receiving neoadjuvant chemotherapy, early reductions in ctDNA variant allele fraction were associated with subsequent treatment response, and patients who achieved a pathological complete response showed substantially greater reductions in ctDNA than those with partial or no response [20]. More recent longitudinal work using ovarian cancer cell-free DNA has similarly shown that resistant clones already present at diagnosis undergo selective expansion during treatment, and that this expansion is trackable in serial plasma samples [21].

Taken together, these studies demonstrate that serial liquid biopsy can provide clinically relevant lead time over conventional assessment in selected settings, establishing a biological and temporal foundation for the development of future forecasting models.

But are these actually AI predictions?

The distinction becomes particularly important when evaluating studies described as AI-based prediction. A serial biomarker study showing that ctDNA changes before radiographic progression establishes temporal association, while an association between ctDNA burden and survival provides prognostic information. Neither necessarily demonstrates that a model can forecast a future event for an individual patient, and even a machine-learning classifier should not automatically be regarded as a forecasting system solely because it achieves high discrimination [25,26]. A useful hierarchy therefore separates five levels of evidence. Serial biomarker measurement demonstrates biological monitoring; statistical association establishes prognostic or predictive relationships; machine-learning classification may improve discrimination for a predefined endpoint; and longitudinal AI modelling explicitly incorporates repeated observations to estimate a subsequent event [17]. Stronger evidence comes from prospective validation in which predictions are generated before the outcome, while AI-guided intervention represents a still higher threshold in which predictions inform treatment and improve patient outcomes. At the clinical-action level, the DYNAMIC trial provides an example of ctDNA-guided clinical decision-making rather than AI-based prediction: postoperative ctDNA status was used to guide adjuvant chemotherapy decisions in stage II colon cancer, demonstrating how molecular information can be prospectively incorporated into treatment management [23]. This hierarchy is important because longitudinal prediction studies vary substantially in their prediction points, horizons, outcomes, and methodological quality. A recent scoping review of AI models applied to longitudinal cancer data identified considerable variation in these features and a high risk of bias in most included studies, particularly from inappropriate study design and inadequate sample size [24]. Broader systematic reviews of machine-learning prediction models similarly report substantial risk of bias, frequently related to insufficient sample size, overfitting, and other analytical shortcomings [25,26]. Claims of AI-enabled prediction should therefore be judged by temporal design, appropriate validation, and demonstrated clinical value rather than algorithmic sophistication alone (Table 1).

Table 1. Evidence hierarchy for AI-enabled tumour prediction.

Evidence type	What it demonstrates	Does it constitute genuine forecasting?	Example
Serial biomarker measurement	Biological change over time	No	Serial ctDNA measurements showing decreasing or increasing tumour burden during treatment
Statistical association	Prognostic or predictive association	Not necessarily	Association between detectable ctDNA and subsequent recurrence or survival
Machine-learning classifier	Classification of a predefined outcome	Usually not by itself	Classification of patients into high- and low-risk groups using molecular features
Longitudinal AI model	Prediction using temporal trajectories	Potentially	Prediction of future progression or treatment resistance using serial ctDNA measurements
Prospectively validated model	Prediction in independently assessed patients	Stronger evidence	Prediction generated before the clinical outcome and subsequently evaluated in patients
AI-guided intervention	Prediction translated into treatment change	Highest level of clinical evidence	Model-generated prediction used to trigger predefined treatment modification or clinical reassessment

To place these concepts in the context of current evidence, studies across different cancer types have investigated the use of serial liquid-biopsy measurements to monitor tumour evolution, treatment response, and emerging resistance. However, the extent to which these approaches achieve genuine longitudinal prediction varies considerably. Table 2 summarises representative evidence across cancer settings, highlighting the analytes examined, longitudinal sampling strategies, analytical

approaches, and the outcomes investigated. Importantly, the studies summarised in Table 2 represent two related but distinct evidence streams: longitudinal liquid-biopsy studies that track molecular changes over time, and studies that explicitly apply AI or machine-learning methods to longitudinal or multimodal data. The former provide biological and temporal evidence necessary for prediction but should not themselves be interpreted as AI-enabled forecasting.

Table 2. Representative evidence for longitudinal liquid-biopsy monitoring and AI/ML-enabled prediction of tumour evolution and treatment resistance.

Study/cancer setting	Design/population	Longitudinal sampling	Analysis	AI/ML?	Endpoint/prediction horizon	Validation/performance
NSCLC [17]	Phase III trial-derived cohort; 466 patients with metastatic NSCLC receiving chemo-immunotherapy	ctDNA assessed at 5 time points through cycle 3 day 1	ML integration of longitudinal ctDNA metrics	Yes	Overall survival; risk estimated using early on-treatment ctDNA measurements	External validation in OAK cohort; OS HR 3.73 (95% CI 1.83–7.60) for high-risk group
Metastatic colorectal cancer [16]	Longitudinal cohort; 62 patients receiving first-line chemotherapy; 272 plasma samples	Serial sampling during treatment	Targeted NGS and molecular trajectory analysis	No	Radiological progression/PFS; ctDNA progression preceded radiological progression by median 3.3 months	ctDNA clearance associated with PFS; adjusted HR 0.22 (95% CI 0.10–0.46); no external predictive-model validation
EGFR-mutant NSCLC [18,19]	[18]: 117 evaluable patients with acquired EGFR-TKI resistance; [19]: randomised non-	Serial plasma T790M monitoring	ddPCR/co bas-based molecular	No	Emergence of T790M/resistance before clinical or RECIST progression	[18]: T790M detected median 2.2 months before clinical PD in

Study/cancer setting	Design/population	Longitudinal sampling	Analysis	AI/ML?	Endpoint/prediction horizon	Validation/performance
	comparative phase II APPLE trial, 103 patients in reported arms B/C	during EGFR-TKI therapy	monitoring			approximately half of T790M-positive patients; [19]: 17% switched treatment before RECIST PD; 18-month PFS 67.2% in arm B
Early breast cancer [20]	Prospective cohort; 52 patients receiving neoadjuvant chemotherapy	Plasma collected before, during and after chemotherapy	Serial ctDNA/VAF analysis and ROC analysis	No	Treatment response and prognosis; response assessed from ctDNA after 2 treatment cycles	AUC 0.81 (95% CI 0.61–1.00) for local tumour response after 2 cycles; no external validation
High-grade serous ovarian cancer [21]	Prospective multimodal cohort; 24 stage III–IV patients; longitudinal clonal analysis in 18 patients	Serial cfDNA from diagnosis through treatment and recurrence over multiple years	CloneSeq-SV and evolutionary modelling of clone-specific structural variants	No*	Clonal evolution and emergence of treatment-resistant populations; no predefined clinical prediction horizon	Demonstrated selective expansion of pre-existing resistant clones; no external clinical forecasting validation
Study/cancer setting	Design/population	Longitudinal sampling	Analysis	AI/ML?	Endpoint/prediction horizon	Validation/performance
	Phase III trial-derived cohort; 466 patients with metastatic NSCLC receiving chemo-immunotherapy	ctDNA assessed at 5 time points through cycle 3 day 1	ML integration of longitudinal ctDNA metrics	Yes	Overall survival; risk estimated using early on-treatment ctDNA measurements	External validation in OAK cohort; OS HR 3.73 (95% CI 1.83–7.60) for high-risk group

Although Williams et al. [21] employed computational evolutionary modelling through CloneSeq-SV, the study did not evaluate an AI/ML model for prospective patient-level clinical forecasting; it is therefore classified here as longitudinal clonal-evolution evidence rather than AI-enabled prediction.

The prediction and action gap

Even an accurate prediction of treatment resistance does not automatically translate into better patient care. A clinically useful model must provide sufficient lead time for a meaningful intervention and, importantly, the predicted outcome must have an available and effective response. If an algorithm identifies a resistant trajectory but there is no validated treatment strategy for acting on that information, its clinical value remains uncertain. The same applies when changing therapy solely on the basis of a molecular prediction has not been shown to improve outcomes. This creates a critical gap between predicting what a tumour may do and knowing what should be done in response. Bridging this gap will require

prospective clinical trials in which AI-generated predictions are linked to predefined treatment strategies and evaluated not only for predictive accuracy, but also for patient outcomes, safety, and clinical utility [23].

The prediction–action gap can be illustrated by several potential clinical scenarios. A validated forecast of emerging molecular resistance during targeted therapy could prompt confirmatory testing, closer surveillance, or consideration of an alternative treatment before radiological progression becomes evident. In the postoperative setting, a persistent or rising ctDNA trajectory indicating a high probability of recurrence could potentially inform escalation of adjuvant therapy or intensified surveillance, whereas sustained molecular clearance might support investigation of treatment de-escalation strategies. Similarly, a forecast of progression generated from serial ctDNA measurements during systemic therapy could identify patients who require earlier imaging or clinical reassessment. Importantly, these actions should

not be triggered solely because a model predicts an adverse outcome; prospective trials must demonstrate that acting on such predictions improves clinically

meaningful outcomes compared with standard management.

Why Prediction Remains Difficult

Biological limitations

The predictive value of liquid biopsy is constrained first by the biology of what reaches the circulation. Tumour shedding is highly variable between patients and across disease stages, and some tumours release very little ctDNA, limiting the amount of tumour-derived information available for analysis. Even when ctDNA is detectable, plasma represents a composite signal from multiple tumour sites and may therefore underrepresent spatially restricted clones. A resistant population emerging in one metastatic site, for example, may not contribute sufficiently to the circulating pool to become detectable at the same time as it becomes biologically important. This creates a fundamental problem for longitudinal prediction because the absence of a circulating signal does not necessarily indicate the absence of an evolving resistant clone. In addition, circulating DNA is not exclusively tumour derived; cell-free DNA released by normal tissues contributes background material that can dilute tumour-specific signals and complicate their accurate measurement. These biological constraints mean that an AI model may identify patterns with considerable technical precision while still providing an incomplete representation of the tumour from which those patterns originate.

The magnitude of these biological limitations varies substantially across clinical contexts. ctDNA shedding differs among tumour types and generally increases with greater disease burden, meaning that molecular trajectories derived from advanced metastatic disease may not translate directly to early-stage or minimal residual disease settings, where circulating tumour-derived material can be extremely scarce. Metastatic site also influences detectability because lesions in different anatomical compartments may contribute unequally to the circulating DNA pool. Treatment further alters ctDNA dynamics by changing tumour burden, inducing tumour-cell death, and selecting resistant populations, such that the biological meaning of a rise or fall in ctDNA may depend on when sampling occurs relative to therapy. These context-dependent differences are particularly important for AI models because a trajectory learned in one tumour type, disease stage, metastatic pattern, or treatment setting may not generalise to another. Longitudinal prediction models should therefore account explicitly for tumour type, disease burden, metastatic distribution, treatment exposure, and the clinical context in which ctDNA is

measured [37]. Clonal haematopoiesis represents an additional biological confounder in cfDNA analysis because age-related expansion of haematopoietic cell clones can release somatic variants into plasma that may be mistaken for tumour-derived alterations. This is particularly relevant when variants occur in genes commonly altered in both haematological clones and solid tumours. Without appropriate discrimination, such signals may generate false-positive ctDNA findings or distort longitudinal molecular trajectories used for AI modelling. Parallel sequencing of matched white blood cells can help distinguish clonal-haematopoiesis-derived variants from tumour-derived mutations and should therefore be considered when constructing and validating longitudinal ctDNA datasets [36].

Analytical limitations

Analytical variability remains an important limitation of longitudinal liquid-biopsy modelling. Sequencing errors can become particularly problematic when tumour-derived DNA is present at low abundance, making small changes in variant frequency difficult to distinguish from technical noise. This is important for AI-based longitudinal analysis because a model may interpret technical variation as evidence of biological change.

Assay differences introduce a further challenge. Differences in extraction methods, sequencing platforms, analytical pipelines, and reporting thresholds can produce measurements that are not directly comparable across laboratories or across time. Pre-analytical factors, including blood collection, processing, and storage, can also influence the quantity and quality of circulating DNA and therefore affect downstream measurements.

Assay design further influences the sensitivity and consistency of longitudinal ctDNA measurements. Tumour-informed approaches use prior tumour sequencing to identify patient-specific variants that are subsequently tracked in plasma, enabling highly targeted detection and potentially greater sensitivity for low-level disease, although they require tumour tissue and prior molecular profiling. In contrast, tumour-agnostic approaches analyse plasma without requiring prior tumour profiling, typically using predefined genomic panels or broader molecular features such as methylation and fragmentomic patterns. These approaches may be more readily implemented when tumour tissue is unavailable and may capture newly

emerging alterations, but their sensitivity and specificity depend on assay design and the biological signal being measured. Consequently, longitudinal changes generated by different assay strategies may not be directly comparable, and apparent molecular trajectories may partly reflect differences in analytical sensitivity, feature selection, or detection thresholds rather than tumour evolution itself. Assay type should therefore be explicitly considered when developing and validating AI models based on serial liquid-biopsy measurements [34,35].

These analytical issues become more consequential when serial measurements are used to construct a tumour trajectory. A change between two sampling points is meaningful only if it can be distinguished from expected analytical variation. Lack of standardised procedures and thresholds therefore limits both reproducibility and the transferability of longitudinal models across clinical settings. Apparent temporal changes in ctDNA should consequently be interpreted in the context of the assay strategy used, particularly when comparing or combining longitudinal datasets generated using different approaches.

The key issue is that AI cannot correct for unreliable input data simply by using a more sophisticated algorithm. If technical variation is incorporated into the training data, the model may learn patterns that do not represent genuine tumour evolution. Standardised sample handling, validated assays, and consistent analytical procedures are consequently necessary foundations for reliable AI-enabled prediction.

AI-specific limitations

AI-based longitudinal modelling introduces additional sources of uncertainty beyond those associated with liquid-biopsy measurement. One of the most important is the limited size of many available datasets. Longitudinal studies require repeated samples from the same patients, which makes them more difficult and expensive to assemble than conventional cross-sectional datasets. When the number of patients is small relative to the number of molecular features, models can become overly complex and overfit the development dataset, producing apparently strong performance that does not generalise to new patients [27].

Data leakage and temporal bias are major concerns in longitudinal prediction and can arise in several ways. Randomly splitting individual observations rather than patients can place serial samples from the same patient in both training and testing datasets, allowing the model to exploit patient-specific information and producing overly optimistic

performance. Leakage can also occur when biomarker, imaging, treatment, or clinical information collected after the defined prediction point is incorporated into model development. Immortal-time bias may arise when classification or exposure definitions depend on patients remaining event-free long enough to contribute subsequent measurements, while inappropriate imputation can introduce future information if missing values at an earlier time point are estimated using measurements obtained later in the disease course. Class imbalance can further produce misleading estimates of accuracy when clinically important outcomes, such as treatment resistance, occur in relatively few patients. Longitudinal prediction studies should therefore use patient-level data partitioning, temporally appropriate preprocessing and imputation, clearly defined prediction points and horizons, and validation procedures that preserve the chronological ordering of information.

Longitudinal datasets also commonly contain missing or irregularly timed observations. Patients may miss sampling points, discontinue treatment, or undergo testing at different intervals. These patterns are unlikely to be completely random and may themselves be related to disease progression, creating additional challenges for model development.

Finally, strong performance in a single institution does not establish clinical reliability. Poor external validation, dataset shift, and institution-specific biases can reduce performance when models are applied to patients whose tumour types, treatment patterns, assays, or clinical characteristics differ from those in the development cohort [30]. Lack of interpretability further limits clinical adoption when clinicians cannot understand which longitudinal features are driving a prediction. Thus, the main challenge is not simply developing accurate algorithms, but demonstrating that their predictions remain reliable, interpretable, and clinically meaningful beyond the dataset in which they were trained.

The biggest problem: prediction without clinical validation

A model can achieve an impressive area under the receiver operating characteristic curve (AUROC) and still have limited clinical value. Technical performance describes how well a model discriminates between patients who do and do not experience a specified outcome, but this does not establish whether its predictions are sufficiently accurate, timely, or useful to change clinical management. In longitudinal liquid-biopsy research, this distinction is particularly important because a model may identify patients at

high risk of progression without demonstrating that acting on that prediction improves their outcome.

Clinical utility depends on factors that are not captured by discrimination alone. Predictions should be adequately calibrated, meaning that estimated risks should correspond reasonably to the risks actually observed [28]. The model should also provide sufficient lead time for clinicians to intervene and should perform reliably in patients outside the development cohort. Importantly, there must be a clinically justified action associated with the prediction. Identifying an increasing probability of treatment resistance is of limited practical value if there is no validated treatment strategy that can safely be initiated on the basis of that information.

The distinction is therefore between predicting an outcome and improving a decision. External and

prospective validation can establish whether a model continues to perform in real clinical settings [29], but even good predictive performance does not by itself demonstrate patient benefit. The strongest evidence would come from prospective studies in which model predictions are available before the clinical event and are linked to predefined treatment strategies, with outcomes compared against standard care [23].

For AI-enabled liquid biopsy to move beyond research, evaluation must therefore extend from AUROC and other measures of technical performance to calibration, clinical utility, decision consequences, and patient outcomes [31]. This is the point at which a promising prediction model becomes a clinically relevant tool rather than simply a high-performing algorithm.

From Prediction to Clinical Decision-Making

The potential clinical pathway for AI-enabled liquid biopsy is straightforward in concept but remains largely investigational: serial liquid biopsy, AI analysis, predicted evolutionary trajectory, estimated resistance risk, clinical decision, and repeat sampling [33]. In this model, repeated molecular measurements would provide an evolving picture of tumour behaviour,

while AI would estimate the probability of clinically relevant changes before they become evident through conventional assessment. The resulting prediction could then be considered alongside imaging, clinical findings, and treatment history before a management decision is made. This represents a proposed framework rather than an established clinical pathway (Figure 1).




Figure 1. Conceptual framework for AI-enabled longitudinal oncology.

Serial liquid-biopsy sampling generates repeated molecular measurements that can be integrated into patient-specific molecular trajectories. AI-based temporal modelling uses these longitudinal data to forecast tumour evolution and estimate the risk of emerging treatment resistance. These predictions may inform treatment adaptation, after which subsequent liquid-biopsy measurements provide new information on tumour dynamics and treatment response. The newly acquired longitudinal data can then be incorporated to update and refine the predictive model, creating a continuous learning cycle. This framework represents a proposed conceptual pathway and remains investigational pending prospective validation and demonstration of clinical utility.

Translation into practice would require prospective validation and clinical trials demonstrating that acting on model predictions improves outcomes compared with standard care [23]. Interpretability is also important because clinicians must be able to assess why a model has identified a patient as being at increased risk, rather than treating its output as an unquestionable result. Clinician oversight should

therefore remain central, particularly when predictions are uncertain or conflicting with other clinical evidence.

Practical considerations may be equally important. The cost of repeated liquid-biopsy testing, sequencing, and computational analysis could limit access, particularly in healthcare systems with constrained resources. Differences in laboratory capacity and availability of advanced molecular testing may also create inequalities in implementation. In addition, models intended for clinical use would need to meet relevant regulatory and quality requirements and demonstrate reliable performance across different patient populations and healthcare settings [29].

Finally, a useful system would require patient-specific decision thresholds. The level of predicted resistance risk that justifies changing treatment may differ according to cancer type, available therapies, treatment toxicity, and the consequences of delaying a more effective option. The goal should therefore not be to replace clinical judgement with an algorithm, but to provide timely longitudinal evidence that can improve clinical decision-making [31].

Future Research Agenda and Proposed Framework

The next stage of AI-enabled liquid-biopsy research should focus on establishing whether longitudinal molecular information can generate reliable and clinically useful forecasts of tumour behaviour. To structure this progression, we propose

the Liquid Biopsy-to-Forecast Continuum, a six-stage framework comprising detection, monitoring, biological characterisation, forecasting, clinical action, and iterative learning.

Table 3. The Liquid Biopsy-to-Forecast Continuum: a proposed six-stage evidentiary framework.

Stage	Objective
1 – Detect	Determine whether tumour-derived material can be reliably identified in the circulation.
2 – Monitor	Assess how the molecular signal changes across serial sampling points.
3 – Characterise	Identify emerging clones, resistance-associated alterations, or other biological changes that may explain the observed trajectory.
4 – Forecast	Use longitudinal molecular and clinical information to estimate the probability of a defined future event within a clinically meaningful time horizon.
5 – Act	Determine whether the prediction is sufficiently reliable and actionable to support a change in treatment.
6 – Learn	Continue longitudinal sampling to determine whether the predicted and observed trajectories agree and, where appropriate, use these data to improve subsequent models.

The Liquid Biopsy-to-Forecast Continuum was developed through conceptual synthesis of the evidence reviewed in this article rather than through formal consensus or empirical validation. Its six stages reflect the progressive evidentiary steps identified across the literature, beginning with reliable detection of tumour-derived material, followed by longitudinal monitoring and biological characterisation, and

advancing to forecasting, clinical action, and iterative learning. The framework can be applied prospectively when designing clinical studies by requiring investigators to specify the stage being evaluated and the corresponding evidence needed for progression to the next stage. For example, forecasting studies should predefine the prediction point, prediction horizon, clinical endpoint, and validation strategy, whereas

studies evaluating the action stage should additionally determine whether model-generated predictions can support predefined treatment decisions and improve clinically meaningful outcomes. The framework therefore provides a structured approach for distinguishing the maturity of different AI-enabled liquid-biopsy applications and identifying the methodological and clinical evidence still required for translation. The proposed continuum is not intended to replace established frameworks for biomarker development, clinical prediction modelling, or implementation. Rather, it addresses a specific translational gap at their intersection by organising the progression from serial liquid-biopsy measurement to prospective, AI-enabled forecasting and ultimately to clinical action. Unlike conventional biomarker frameworks that primarily establish analytical and clinical validity, or prediction-model frameworks that focus on model development, validation, and reporting, the continuum explicitly links temporally resolved molecular monitoring with biological interpretation, future-event forecasting, predefined clinical action, and iterative learning. Its principal contribution is therefore to provide a liquid-biopsy-specific structure for determining when longitudinal molecular information has progressed from detecting or monitoring tumour change to generating and acting on clinically meaningful predictions.

Several priorities follow from this framework. Future studies should use prospective longitudinal cohorts with standardised sampling schedules and clearly defined prediction points and horizons [27]. Multimodal datasets integrating liquid-biopsy measurements with imaging, treatment exposure, and clinical characteristics may provide a more complete representation of tumour state, but these models require rigorous external and temporal validation [29,30]. Studies should also emphasise interpretable models, calibrated risk estimates, and clinically

meaningful prediction thresholds rather than relying on discrimination measures alone [28].

Ultimately, prospective interventional trials are needed to determine whether acting on AI-generated predictions improves patient outcomes, following the precedent set by ctDNA-guided adjuvant chemotherapy trials in colon cancer [23]. Research should also assess cost, accessibility, workflow requirements, and other health-economic and implementation outcomes [31]. The future of AI-enabled liquid biopsy should therefore not be defined by whether an algorithm can classify patients, but by whether longitudinal models can reliably forecast clinically meaningful tumour transitions early enough to change clinical outcomes.

Limitations of this review

This review has several limitations. First, it was conducted as a narrative critical review rather than a systematic review; therefore, the literature was purposively selected and the possibility of selection bias cannot be excluded. Second, no formal quantitative synthesis or meta-analysis was undertaken, reflecting substantial heterogeneity across studies in cancer type, disease stage, treatment setting, liquid-biopsy analyte, sampling schedule, analytical platform, modelling approach, and clinical endpoint. This heterogeneity also limits direct comparison of predictive performance across studies. Third, the evidence base for AI-enabled longitudinal forecasting remains relatively limited, with many studies focusing on molecular monitoring, prognostic association, or early detection rather than prospectively validated patient-level forecasting. Finally, AI and liquid-biopsy technologies are evolving rapidly, and emerging studies may alter the evidence landscape as new assays, longitudinal datasets, modelling approaches, and prospective validation studies become available. These limitations should be considered when interpreting the conclusions and proposed framework presented in this review.

Conclusion

Can liquid biopsy predict the future of a tumour? Current evidence suggests that this capability remains investigational. Serial liquid-biopsy measurements can reveal tumour burden, clonal evolution, treatment response, and emerging resistance before conventional clinical progression in selected settings. However, prospectively validated AI models that use longitudinal liquid-biopsy trajectories to forecast patient-specific tumour evolution or treatment resistance remain limited. Advancing the field will require standardised assays and sampling, adequately powered longitudinal cohorts, rigorous temporal and external validation, calibrated and interpretable

models, and prospective trials demonstrating that model-informed treatment decisions improve patient outcomes. AI-enabled liquid biopsy therefore provides a credible pathway towards anticipatory oncology, but its clinical value will ultimately depend on demonstrating reliable forecasting and actionable benefit beyond molecular monitoring.

Acknowledgements

Funding: The authors received no financial support or funding for the preparation of this manuscript.

Conflicts of interest: The authors declare that they have no conflicts of interest.

Author contributions: All authors contributed to the conception and design of the manuscript, literature review, interpretation of the available evidence, drafting and critical revision of the manuscript, and approved the final version for publication.

Data availability: Not applicable. This article is a narrative review and does not involve the generation or analysis of new datasets.

Ethics approval: Not applicable. This manuscript is a narrative review based on previously published

literature and does not involve human participants, animals, or primary data collection.

Disclosure of AI Use: During the preparation of this manuscript, the authors used ChatGPT (OpenAI) for language editing, refinement of wording, and improvement of readability. The use of the tool was limited to these non-intellectual editorial tasks. All scientific content, interpretation of the literature, conclusions, and final manuscript were reviewed and approved by the authors, who take full responsibility for the accuracy, originality, and integrity of the work.

References

1. Nowell PC. The clonal evolution of tumor cell populations. *Science*. 1976;194(4260):23-8. doi:[10.1126/science.959840](https://doi.org/10.1126/science.959840)
2. McGranahan N, Swanton C. Clonal heterogeneity and tumor evolution: past, present, and the future. *Cell*. 2017;168(4):613-28. doi:[10.1016/j.cell.2017.01.018](https://doi.org/10.1016/j.cell.2017.01.018)
3. Dagogo-Jack I, Shaw AT. Tumour heterogeneity and resistance to cancer therapies. *Nat Rev Clin Oncol*. 2018;15(2):81-94. doi:[10.1038/nrclinonc.2017.166](https://doi.org/10.1038/nrclinonc.2017.166)
4. Siravegna G, Marsoni S, Siena S, Bardelli A. Integrating liquid biopsies into the management of cancer. *Nat Rev Clin Oncol*. 2017;14(9):531-48. doi:[10.1038/nrclinonc.2017.14](https://doi.org/10.1038/nrclinonc.2017.14)
5. Wan JCM, Massie C, Garcia-Corbacho J, Mouliere F, Brenton JD, Caldas C, et al. Liquid biopsies come of age: towards implementation of circulating tumour DNA. *Nat Rev Cancer*. 2017;17(4):223-38. doi:[10.1038/nrc.2017.7](https://doi.org/10.1038/nrc.2017.7)
6. Bettegowda C, Sausen M, Leary RJ, Kinde I, Wang Y, Agrawal N, et al. Detection of circulating tumor DNA in early- and late-stage human malignancies. *Sci Transl Med*. 2014;6(224):224ra24.
7. Alix-Panabières C, Pantel K. Clinical applications of circulating tumor cells and circulating tumor DNA as liquid biopsy. *Cancer Discov*. 2016;6(5):479-91. doi:[10.1158/2159-8290.CD-15-1483](https://doi.org/10.1158/2159-8290.CD-15-1483)
8. Aceto N, Bardia A, Miyamoto DT, Donaldson MC, Wittner BS, Spencer JA, et al. Circulating tumor cell clusters are oligoclonal precursors of breast cancer metastasis. *Cell*. 2014;158(5):1110-22. doi:<https://doi.org/10.1016/j.cell.2014.07.013>
9. Gremmelspacher D, Gawron J, Szczerba BM, Jahn K, Castro-Giner F, Kuipers J, et al. Phylogenetic inference reveals clonal heterogeneity in circulating tumor cell clusters. *Nat Genet*. 2025;57(6):1357-61. doi:<https://doi.org/10.1038/s41588-025-02205-2>
10. Kalluri R, LeBleu VS. The biology, function, and biomedical applications of exosomes. *Science*. 2020;367(6478):eaau6977. doi:[10.1126/science.aau6977](https://doi.org/10.1126/science.aau6977)
11. Zhang Y, Sun B, Yu Y, Lu J, Lou Y, Qian F, et al. Multimodal fusion of liquid biopsy and CT enhances differential diagnosis of early-stage lung adenocarcinoma. *NPJ Precis Oncol*. 2024;8:47. doi:<https://doi.org/10.1038/s41698-024-00551-8>
12. Cohen JD, Li L, Wang Y, Thoburn C, Afsari B, Danilova L, et al. Detection and localization of surgically resectable cancers with a multi-analyte blood test. *Science*. 2018;359(6378):926-30. doi:<https://doi.org/10.1126/science.aar3247>
13. Rajkomar A, Dean J, Kohane I. Machine learning in medicine. *N Engl J Med*. 2019;380(14):1347-58. doi:[10.1056/NEJMr1814259](https://doi.org/10.1056/NEJMr1814259)

14. Hochreiter S, Schmidhuber J. Long short-term memory. *Neural Comput.* 1997;9(8):1735-80. doi:[10.1162/neco.1997.9.8.1735](https://doi.org/10.1162/neco.1997.9.8.1735)
15. Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, et al. Attention is all you need. In: *Advances in Neural Information Processing Systems* 30 (NeurIPS 2017). 2017:5998-6008.
16. Kim ST, Lee WS, Lanman RB, Mortimer S, Zill OA, Kim KM, et al. Dynamic changes in longitudinal circulating tumour DNA profile during metastatic colorectal cancer treatment. *Br J Cancer.* 2022;127(4):661-9. <https://doi.org/10.1038/s41416-022-01837-z>
17. Assaf ZJ, Zou W, Fine AD, Socinski MA, Young A, Lloyd A, et al. A longitudinal circulating tumor DNA-based model associated with survival in metastatic non-small-cell lung cancer. *Nat Med.* 2023;29(4):859-68. <https://doi.org/10.1038/s41591-023-02226-6>
18. Zheng D, Ye X, Zhang MZ, Sun Y, Wang JY, Ni J, et al. Plasma EGFR T790M ctDNA status is associated with clinical outcome in advanced NSCLC patients with acquired EGFR-TKI resistance. *Sci Rep.* 2016;6:20913. <https://doi.org/10.1038/srep20913>
19. Remon J, Besse B, Ponce Aix S, Callejo A, Al-Rabi K, Bernabe R, et al. Osimertinib treatment based on plasma T790M monitoring in patients with EGFR-mutant non-small-cell lung cancer (NSCLC): EORTC Lung Cancer Group 1613 APPLE phase II randomized clinical trial. *Ann Oncol.* 2023;34(5):468-76. <https://doi.org/10.1016/j.annonc.2023.02.012>
20. Li S, Lai H, Liu J, Liu Y, Jin L, Li Y, et al. Circulating tumor DNA predicts the response and prognosis in patients with early breast cancer receiving neoadjuvant chemotherapy. *JCO Precis Oncol.* 2020;4:244-57. <https://doi.org/10.1200/PO.19.00292>
21. Williams MJ, Vázquez-García I, Tam G, Wu M, Varice N, Havasov E, et al. Tracking clonal evolution during treatment in ovarian cancer using cell-free DNA. *Nature.* 2025;647(8090):757-65. doi:[10.1038/s41586-025-09580-0](https://doi.org/10.1038/s41586-025-09580-0)
22. Tie J, Wang Y, Tomasetti C, Li L, Springer S, Kinde I, et al. Circulating tumor DNA analysis detects minimal residual disease and predicts recurrence in patients with stage II colon cancer. *Sci Transl Med.* 2016;8(346):346ra92. <https://doi.org/10.1126/scitranslmed.aaf6219>
23. Tie J, Cohen JD, Lahouel K, Lo SN, Wang Y, Kosmider S, et al. Circulating tumor DNA analysis guiding adjuvant therapy in stage II colon cancer. *N Engl J Med.* 2022;386(24):2261-72. <https://doi.org/10.1056/NEJMoa2200075>
24. Moglia V, Johnson O, Cook G, de Kamps M, Smith L. Artificial intelligence methods applied to longitudinal data from electronic health records for prediction of cancer: a scoping review. *BMC Med Res Methodol.* 2025;25:24. <https://doi.org/10.1186/s12874-025-02473-w>
25. Andaur Navarro CL, Damen JAA, Takada T, Nijman SWJ, Dhiman P, Ma J, et al. Risk of bias in studies on prediction models developed using supervised machine learning techniques: systematic review. *BMJ.* 2021;375:n2281. <https://doi.org/10.1136/bmj.n2281>
26. Dhiman P, Ma J, Andaur Navarro CL, Speich B, Bullock G, Damen JAA, et al. Risk of bias of prognostic models developed using machine learning: a systematic review in oncology. *Diagn Progn Res.* 2022;6:13. <https://doi.org/10.1186/s41512-022-00126-w>
27. Riley RD, Ensor J, Snell KIE, Harrell FE Jr, Martin GP, Reitsma JB, et al. Calculating the sample size required for developing a clinical prediction model. *BMJ.* 2020;368:m441. <https://doi.org/10.1136/bmj.m441>
28. Van Calster B, McLernon DJ, van Smeden M, Wynants L, Steyerberg EW. Calibration: the Achilles heel of predictive analytics. *BMC Med.* 2019;17:230. <https://doi.org/10.1186/s12916-019-1466-7>
29. Collins GS, Moons KGM, Dhiman P, Riley RD, Beam AL, Van Calster B, et al. TRIPOD+AI statement: updated guidance for reporting clinical prediction models that use regression or machine learning methods. *BMJ.* 2024;385:e078378. <https://doi.org/10.1136/bmj-2023-078378>

30. Futoma J, Simons M, Panch T, Doshi-Velez F, Celi LA. The myth of generalisability in clinical research and machine learning in health care. *Lancet Digit Health*. 2020;2(9):e489-92. [https://doi.org/10.1016/S2589-7500\(20\)30186-2](https://doi.org/10.1016/S2589-7500(20)30186-2)
31. Topol EJ. High-performance medicine: the convergence of human and artificial intelligence. *Nat Med*. 2019;25(1):44-56. doi:[10.1038/s41591-018-0300-7](https://doi.org/10.1038/s41591-018-0300-7)
32. Oluwadare OE, Frankpeace MS, Oluwaniran OM, Anatuanya JI, Onwuemelem LA. CRISPR Functional Genomics in Precision Oncology: Integrating Single-Cell Multi-Omics for Cancer Vulnerability Discovery. *Oncology, Nuclear Medicine and Transplantology*. 2026;2(3):onmt027. doi:[10.63946/onmt/19194](https://doi.org/10.63946/onmt/19194)
33. Oluwadare OE, Adeoba MI, Akor JT, Yusuff TA, Milimo P. The Rise of Smart Hospitals: Biomedical Engineering Innovations in Automation, Monitoring, and Digital Healthcare. *Journal of Medicine and Health Research*. 2026;11(2):121-38. doi:[10.56557/jomahr/2026/v11i210794](https://doi.org/10.56557/jomahr/2026/v11i210794)
34. Malla M, Loree JM, Kasi PM, Parikh AR. Using circulating tumor DNA in colorectal cancer: current and evolving practices. *J Clin Oncol*. 2022;40(24):2846-57. doi:[10.1200/JCO.21.02615](https://doi.org/10.1200/JCO.21.02615)
35. Dasari A, Morris VK, Allegra CJ, Atreya C, Benson AB 3rd, Boland P, et al. ctDNA applications and integration in colorectal cancer: an NCI Colon and Rectal–Anal Task Forces whitepaper. *Nat Rev Clin Oncol*. 2020;17(12):757-70. doi:[10.1038/s41571-020-0392-0](https://doi.org/10.1038/s41571-020-0392-0)
36. Razavi P, Li BT, Brown DN, Jung B, Hubbell E, Shen R, et al. High-intensity sequencing reveals the sources of plasma circulating cell-free DNA variants. *Nat Med*. 2019;25(12):1928-37. doi:[10.1038/s41591-019-0652-7](https://doi.org/10.1038/s41591-019-0652-7)
37. Pantel K, Alix-Panabières C. Minimal residual disease as a target for liquid biopsy in patients with solid tumours. *Nat Rev Clin Oncol*. 2025;22:65-77. <https://doi.org/10.1038/s41571-024-00967-y>