

Integrating Circulating Tumor DNA Dynamics and Molecular Biomarkers for Longitudinal Precision Oncology Monitoring Through Data Driven Cancer Surveillance Models

Rukayat Olayemi Oni¹

¹Department of Biology, Georgia State University, Atlanta, Georgia

Publication Date: 2023/01/28

Abstract

Longitudinal precision oncology requires surveillance methods capable of distinguishing clinically meaningful molecular changes from biological variability, sequencing noise, and transient fluctuations in circulating tumor DNA. Conventional circulating tumor DNA surveillance commonly relies on fixed variant allele frequency thresholds, isolated molecular measurements, or predictive models that inadequately represent temporal dependencies among tumor-derived signals. This study develops a data-driven cancer surveillance framework integrating longitudinal circulating tumor DNA dynamics with molecular biomarkers for early detection of disease progression, treatment response, molecular residual disease, and recurrence. A novel *Longitudinal Multimodal Cancer Surveillance Network (LMCS-Net)* is proposed to jointly model circulating tumor DNA variant allele frequency, mutant molecules per milliliter, circulating tumor DNA fraction, fragmentomic characteristics, copy-number alterations, methylation signatures, somatic mutation burden, selected protein biomarkers, treatment exposure, and sequential sampling intervals. LMCS-Net combines adaptive molecular-noise suppression, temporal attention, gated biomarker fusion, patient-specific baseline normalization, Bayesian uncertainty calibration, and dynamic risk-state estimation. For patient (i) at monitoring time (t), the multimodal observation vector ($\mathbf{x}_{i,t}$) is transformed into a latent molecular state ($\mathbf{h}_{i,t}$), from which a continuous cancer-surveillance risk score ($R_{i,t} \in [0,1]$) and molecular progression probability are estimated. Temporal change detection further evaluates ($\Delta R_{i,t} = R_{i,t} - R_{i,t-1}$), enabling clinically significant trajectory changes to be distinguished from isolated biomarker excursions.

The proposed model is benchmarked against fixed-threshold ctDNA classification, logistic regression, random forest, XGBoost, long short-term memory networks, gated recurrent units, conventional transformers, and Bayesian change-point detection. Performance evaluation incorporates area under the receiver operating characteristic curve, area under the precision-recall curve, sensitivity, specificity, F1-score, Brier score, calibration error, concordance index, false-alert rate, and median lead time for detecting molecular progression before conventional clinical confirmation. Comparative analyses are designed to determine whether multimodal temporal fusion improves sensitivity for low-abundance ctDNA while maintaining robust specificity and calibration. Results are visualized through receiver operating characteristic and precision-recall curves, longitudinal ctDNA-risk trajectories, calibration plots, lead-time distributions, biomarker contribution graphs, patient-level molecular state maps, and comparative algorithm-performance graphs. Ablation and sensitivity analyses quantify the contributions of temporal modeling, fragmentomic features, molecular biomarker fusion, uncertainty calibration, and individualized baselines. The proposed framework establishes an interpretable computational architecture for converting repeated molecular measurements into dynamic cancer-surveillance states and provides a foundation for personalized, data-driven oncology monitoring across treatment and post-treatment follow-up.

Keywords: Circulating Tumor DNA; Molecular Biomarkers; Longitudinal Monitoring; Precision Oncology; Cancer Surveillance.

I. INTRODUCTION

➤ *Background to Longitudinal Precision Oncology Monitoring*

Longitudinal precision oncology is moving cancer management from episodic assessment toward continuous characterization of disease state across diagnosis, treatment, remission, and recurrence. Conventional tissue biopsy remains essential for diagnosis and genomic characterization, yet a single specimen captures only one anatomical site and one moment in an evolving disease process. Tumors undergo clonal selection under therapeutic pressure and can acquire resistance-associated alterations absent from the pretreatment biopsy. Serial tissue sampling is also constrained by invasiveness, cost, and limited accessibility of some lesions. Circulating tumor DNA (ctDNA) provides a minimally invasive route for repeated molecular assessment because tumor-derived fragments in plasma can reflect genomic alterations from multiple disease sites. This creates a basis for monitoring molecular response, emerging resistance, minimal residual disease, and recurrence between radiological assessments. Wan et al. (2017) emphasized that the clinical value of ctDNA extends beyond mutation detection to repeated quantitative disease monitoring, while Corcoran and Chabner (2018) described cell-free DNA analysis as a means of tracking treatment-sensitive and treatment-resistant tumor populations over time. The analytical problem is therefore how repeated molecular measurements should be transformed into an interpretable patient-specific trajectory. Longitudinal surveillance must account for baseline tumor burden, treatment exposure, sampling interval, assay variability, molecular heterogeneity, and the possibility that meaningful change is distributed across correlated biomarkers rather than a single ctDNA measurement. This study addresses that requirement through the Longitudinal Multimodal Cancer Surveillance Network (LMCS-Net), which represents each patient as a sequence of time-indexed molecular states. The framework integrates ctDNA variant allele frequency, mutant molecules per milliliter, tumor fraction, fragmentomic characteristics, methylation signals, copy-number alterations, somatic mutation information, selected protein biomarkers, and treatment context. Temporal attention and patient-specific normalization are intended to distinguish persistent molecular progression from transient fluctuations, whereas calibrated risk estimation converts the evolving profile into a dynamic surveillance score. This formulation motivates later comparison with threshold rules, logistic regression, random forest, XGBoost, recurrent neural networks, transformers, and Bayesian change-point methods using discrimination, calibration, false-alert rate, and progression lead time.

➤ *Circulating Tumor DNA as a Dynamic Indicator of Tumor Burden*

Circulating tumor DNA constitutes the tumor-derived fraction of circulating cell-free DNA and can function as a dynamic indicator of disease burden because its concentration and molecular composition change with

tumor growth, cell death, treatment, and biological clearance. Unlike anatomical imaging, which records macroscopic changes at discrete intervals, ctDNA can reveal molecular changes before measurable alterations in lesion dimensions become apparent. Quantitative descriptors include variant allele frequency, absolute mutant molecules per milliliter of plasma, estimated ctDNA fraction, number of detectable tumor-specific variants, and copy-number burden. Diehl et al. (2008) demonstrated that circulating mutant DNA could track tumor dynamics after surgery in colorectal cancer and that persistent or re-emerging mutant DNA was associated with residual or recurrent disease. Dawson et al. (2013) subsequently showed in metastatic breast cancer that serial ctDNA measurements exhibited a broader dynamic range and closer association with changes in tumor burden than CA 15-3 or circulating tumor cells, while providing an earlier indication of treatment response in a subset of patients. These findings support treating ctDNA as a time-dependent signal rather than a binary positive-or-negative marker. The same patient may exhibit substantial variation in ctDNA release according to tumor vascularity, metastatic site, treatment mechanism, clonal composition, and background cell-free DNA. A decrease in variant allele frequency after therapy may indicate molecular response, whereas a sustained increase, reappearance of suppressed variants, or emergence of resistance-associated clones may indicate progression. LMCS-Net is therefore designed to model both the level and direction of change in ctDNA-related variables. For patient (i) at time (t), the surveillance state is informed by the current molecular vector and its deviation from prior measurements, enabling persistent trajectory shifts to be distinguished from isolated excursions. This representation supports outputs such as molecular response probability, progression risk, recurrence risk, and estimated lead time relative to conventional confirmation. The later results evaluate whether integrating several ctDNA metrics improves receiver operating characteristic and precision-recall performance, sensitivity at clinically relevant specificity, concordance, calibration, and false-alert control compared with fixed-threshold or single-time-point methods.

➤ *Molecular Biomarkers in Cancer Surveillance and Treatment Monitoring*

Molecular surveillance becomes more informative when ctDNA mutation measurements are interpreted alongside orthogonal features that capture other dimensions of tumor biology. Somatic single-nucleotide variants and small insertions or deletions provide high sequence specificity, but their detectability declines when tumor fraction is low or when the monitored clone is not representative of dominant disease. Fragmentomic features offer complementary information because tumor-associated cell-free DNA exhibits altered fragment-size distributions and genome-wide fragmentation patterns that reflect nucleosome positioning and chromatin organization. Cristiano et al. (2019) showed that genome-wide fragmentation profiles could distinguish individuals with cancer from those without cancer and contribute information about tissue of origin. Epigenetic biomarkers

provide another signal layer. Liu et al. (2020) demonstrated that targeted methylation signatures in cell-free DNA could detect and localize multiple cancer types with high specificity, illustrating how methylation can capture tumor-associated information without depending on a particular somatic mutation.

For longitudinal precision oncology, the opportunity is to integrate these partially independent molecular channels so that one feature can compensate when another becomes weak or ambiguous. LMCS-Net therefore incorporates ctDNA variant allele frequency, mutant molecules per milliliter, tumor fraction, fragmentomic descriptors, methylation signatures, copy-number alterations, somatic mutation burden, and selected circulating protein biomarkers within a common temporal representation. The objective is not to treat all biomarkers as equally informative, but to learn context-dependent weights through gated multimodal fusion. A rising copy-number abnormality burden accompanied by increasing ctDNA fraction and a concordant fragmentomic shift, for example, should carry stronger evidence of progression than a small isolated increase in one variant allele frequency. Conversely, stable methylation and fragmentomic profiles may reduce confidence assigned to a borderline mutation signal. This design connects directly to the later biomarker-contribution graphs, feature-attribution analyses, and ablation experiments that test the incremental value of each modality. The surveillance score is therefore conceived as a composite molecular state estimate in which temporal direction, cross-biomarker agreement, and uncertainty jointly inform treatment-response and recurrence monitoring.

➤ *Limitations of Conventional Threshold-Based ctDNA Interpretation*

Conventional ctDNA interpretation frequently reduces a continuous and noisy molecular signal to a binary call based on whether variant allele frequency, mutant molecule count, or another measurement exceeds a predefined limit of detection or clinical threshold. Such rules are operationally simple, but they can become unstable near the detection boundary, particularly in minimal residual disease and early recurrence where only a small number of tumor-derived molecules may be present. Sequencing artifacts, polymerase errors, stochastic sampling, low plasma input, and nonuniform molecular recovery can create false signals or obscure true low-frequency variants. Newman et al. (2016) demonstrated that integrated digital error suppression combining molecular barcoding with background-error modeling markedly improved ctDNA detection, illustrating that observed low-frequency variants cannot be interpreted reliably without explicit treatment of technical noise. A fixed threshold also ignores the trajectory of a patient whose values remain below a population cutoff but rise consistently over several sampling points. Biological confounding further limits threshold-based interpretation. Plasma cell-free DNA contains material derived from several tissues, and somatic variants originating from clonal hematopoiesis can be mistaken for tumor-derived alterations. Razavi et al. (2019) showed through high-

intensity sequencing with matched white-blood-cell analysis that a substantial proportion of plasma variants, including many detected in patients with cancer, can originate from clonal hematopoiesis rather than the tumor. Consequently, a single positive molecular measurement does not necessarily represent progression, and a negative measurement does not necessarily establish absence of disease. LMCS-Net addresses these limitations by replacing a static decision boundary with patient-specific baseline normalization, adaptive molecular-noise suppression, temporal evidence accumulation, and Bayesian uncertainty calibration. Instead of asking only whether a value exceeds a fixed cutoff, the model evaluates whether the joint molecular state has changed sufficiently, consistently, and confidently relative to the patient's previous trajectory. This distinction underlies the later comparison with conventional thresholding in terms of low-abundance detection sensitivity, specificity, false-alert rate, calibration error, and lead time, thereby testing whether dynamic evidence fusion provides a more reliable surveillance signal.

➤ *Data-Driven Modeling of Longitudinal Cancer Progression*

Data-driven longitudinal modeling provides a mathematical route for converting serial biomarker observations into continuously updated estimates of disease state and future risk. Traditional prognostic systems commonly rely on baseline covariates or single landmark measurements, whereas cancer progression is a temporally evolving process in which the predictive meaning of a biomarker depends on both its current value and its preceding trajectory. Kurtz et al. (2019) introduced dynamic risk profiling that integrates serial tumor biomarkers to update individualized outcome probabilities as new observations become available, demonstrating the broader value of treating repeated measurements as sequential evidence. More recently, Assaf et al. (2023) evaluated longitudinal ctDNA across multiple treatment time points in metastatic non-small-cell lung cancer and used machine learning to jointly model numerous ctDNA metrics for survival prediction. Their findings showed that the prognostic content of ctDNA depends on metric type and sampling time and that combinations of longitudinal features can support risk stratification beyond isolated measurements.

LMCS-Net extends this principle from longitudinal ctDNA summarization to multimodal molecular surveillance. Each observation is represented as a time-indexed vector containing ctDNA abundance measures, mutation characteristics, fragmentomic features, methylation information, copy-number signals, selected proteins, treatment exposure, and the interval since the previous sample. Temporal attention learns which prior observations are most relevant to the current state, while gated fusion determines the contribution of each biomarker modality under changing signal quality. Patient-specific normalization reduces interindividual scale differences, and Bayesian calibration associates each risk estimate with quantified predictive uncertainty. The resulting risk trajectory permits analysis of absolute

surveillance state, rate of change, persistence, and abrupt molecular transitions. This architecture also permits direct benchmarking against logistic regression, random forest, XGBoost, LSTM, GRU, conventional transformer models, and Bayesian change-point detection. The later results are therefore structured around receiver operating characteristic and precision-recall performance, concordance, calibration, low-abundance detection, false-alert burden, and molecular lead time, supplemented by trajectory plots and ablation analyses that determine whether temporal attention and multimodal fusion provide independent performance gains.

➤ *Research Gap in Multimodal Molecular Surveillance*

Despite rapid progress in liquid biopsy and computational oncology, an important methodological gap remains between multimodal biomarker discovery and clinically useful longitudinal surveillance. Multimodal precision-oncology research has shown that combining complementary molecular, imaging, and clinical information can overcome limitations of single-data-source models, yet integration is constrained by heterogeneous scales, missing observations, sparse longitudinal sampling, variable data quality, and limited external validation. Boehm et al. (2022) noted that most precision-oncology approaches remain dominated by individual data modalities and that multimodal integration requires methods capable of representing heterogeneous inputs without losing modality-specific information. In parallel, the expanding cfDNA field encompasses sequence mutations, copy-number changes, methylation, fragment length, genomic coverage, nucleosome-related patterns, and other fragmentomic characteristics. Bruhm et al. (2025) emphasized that these complementary cfDNA features create opportunities for sensitive multi-feature machine-learning approaches but also raise analytical and translational challenges involving signal selection, standardization, and clinical implementation.

The unresolved problem is especially pronounced in surveillance, where the objective differs from one-time cancer detection. A clinically useful system must learn whether a molecular change is meaningful for a particular patient, whether multiple biomarkers support the same interpretation, how much uncertainty surrounds that interpretation, and whether the change persists. Existing approaches often emphasize cross-sectional classification, single ctDNA measures, fixed molecular thresholds, or outcome prediction without jointly modeling individualized baselines, irregular sampling intervals, multimodal interactions, temporal dependencies, and calibrated uncertainty. This study addresses that gap through LMCS-Net, a unified architecture combining adaptive noise suppression, patient-specific normalization, temporal attention, gated multimodal fusion, dynamic risk estimation, and Bayesian confidence calibration. The framework tests whether integrating molecular trajectories yields more reliable progression and recurrence surveillance than fixed-threshold, statistical, ensemble-learning, recurrent, transformer, and change-point baselines. Evaluation therefore includes discrimination, precision-recall behavior, calibration, low-abundance

sensitivity, false-alert rate, progression lead time, robustness to missing or noisy biomarkers, and ablation of model components. These analyses connect the identified literature gap directly to the later results rather than treating multimodal integration as an abstract computational objective.

➤ *Research Aim, Objectives and Research Questions*

The *research aim* is to develop and evaluate a data-driven longitudinal cancer surveillance framework that integrates circulating tumor DNA dynamics with complementary molecular biomarkers to provide patient-specific, temporally resolved, and uncertainty-calibrated monitoring of treatment response, molecular progression, residual disease, and cancer recurrence.

• *Research Objectives*

- ✓ To develop LMCS-Net for joint temporal modeling of ctDNA dynamics, fragmentomics, methylation, copy-number alterations, somatic mutations, selected protein biomarkers, treatment exposure, and sampling intervals.
- ✓ To formulate patient-specific baseline normalization and adaptive molecular-noise suppression methods for improving interpretation of low-abundance ctDNA measurements during longitudinal surveillance.
- ✓ To develop temporal-attention and gated multimodal-fusion mechanisms capable of converting repeated molecular observations into dynamic cancer-surveillance risk states.
- ✓ To quantify the ability of LMCS-Net to detect treatment response, molecular progression, residual disease, and recurrence earlier than conventional threshold-based surveillance approaches.
- ✓ To benchmark LMCS-Net against logistic regression, random forest, XGBoost, LSTM, GRU, conventional transformer, Bayesian change-point detection, and fixed-threshold ctDNA models.
- ✓ To evaluate discrimination, precision-recall performance, sensitivity, specificity, F1-score, concordance, calibration, Brier score, false-alert rate, molecular lead time, robustness, and component-specific contributions through comparative and ablation analyses.

• *Research Questions*

- ✓ How effectively can longitudinal ctDNA dynamics and complementary molecular biomarkers be integrated into a unified patient-specific cancer-surveillance model?
- ✓ To what extent do adaptive noise suppression and individualized baseline normalization improve detection of low-abundance and temporally changing ctDNA signals?
- ✓ Does temporal attention combined with gated multimodal biomarker fusion improve molecular progression and recurrence detection relative to single-marker and static prediction approaches?
- ✓ How much molecular lead time can longitudinal risk modeling provide relative to conventional clinical

confirmation of treatment failure, progression, or recurrence?

- ✓ How does LMCS-Net compare with fixed-threshold, statistical, ensemble-learning, recurrent neural-network, transformer, and Bayesian change-point models across discrimination and calibration metrics?
- ✓ Which molecular modalities and LMCS-Net components contribute most strongly to predictive performance, robustness, clinical interpretability, and uncertainty-calibrated longitudinal surveillance?

➤ *Contributions, Novelty and Significance of the Study*

The principal contribution of this study is the formulation of LMCS-Net as a unified longitudinal molecular-surveillance architecture rather than a conventional one-time ctDNA classifier. Its novelty lies in simultaneously integrating quantitative ctDNA dynamics, mutant molecules per milliliter, circulating tumor fraction, genomic alterations, copy-number changes, fragmentomic characteristics, methylation signatures, selected circulating protein biomarkers, treatment information, and irregular sampling intervals within a patient-specific temporal framework. The proposed architecture introduces adaptive molecular-noise suppression to reduce instability at low tumor fractions, individualized baseline normalization to accommodate interpatient heterogeneity, temporal attention to identify clinically informative historical measurements, gated multimodal fusion to assign context-sensitive contributions to molecular modalities, and Bayesian calibration to distinguish a high predicted risk from a high-confidence predicted risk. The study further contributes a systematic benchmarking framework in which LMCS-Net is compared with fixed thresholds, conventional statistical models, ensemble learners, recurrent neural networks, transformer models, and Bayesian change-point detection under identical surveillance outcomes. Its significance extends beyond classification accuracy because the evaluation includes progression lead time, false-alert burden, calibration reliability, low-abundance sensitivity, robustness to missing or noisy biomarkers, and patient-level interpretability. This permits the study to determine not merely whether a model predicts progression, but whether it can provide temporally earlier, biologically coherent, and quantitatively reliable evidence that could support precision-oncology surveillance.

➤ *Scope of the Review and Structure of the Paper*

The study focuses on computational integration of serial plasma-derived ctDNA measurements and complementary molecular biomarkers for longitudinal precision-oncology monitoring, with specific emphasis on treatment response, molecular residual disease, progression, and recurrence surveillance. The molecular scope includes variant allele frequency, mutant molecules per milliliter, estimated tumor fraction, somatic mutations, copy-number alterations, fragmentomic features, methylation signatures, and selected circulating protein biomarkers, while treatment exposure and sampling intervals are incorporated as temporal contextual variables. The work concentrates on data-driven surveillance modeling rather than wet-laboratory assay

development, specimen-processing optimization, or therapeutic intervention design. Section 1 establishes the clinical and computational problem, research gap, objectives, and contributions. Section 2 critically reviews ctDNA biology, longitudinal kinetics, molecular biomarkers, conventional interpretation methods, temporal machine learning, and multimodal precision-oncology models. Section 3 develops the mathematical and computational formulation of LMCS-Net, including data representation, normalization, temporal attention, multimodal fusion, uncertainty calibration, training, comparator algorithms, and evaluation procedures. Section 4 presents comparative predictive performance, longitudinal molecular trajectories, progression lead-time analysis, multimodal contribution, calibration, low-abundance detection, ablation, robustness, and interpretability results using numerical tables and graphical comparisons. Section 5 synthesizes the principal findings, clinical implications, limitations, validation requirements, and recommendations for future longitudinal molecular-surveillance systems.

II. LITERATURE REVIEW

➤ *Biological and Molecular Characteristics of Circulating Tumor DNA in Cancer Surveillance*

Circulating tumor DNA (ctDNA) is the tumor-derived component of cell-free DNA released into plasma through apoptosis, necrosis, active secretion, and other tumor–host processes. Its abundance is shaped by tumor volume, vascularity, metastatic site, cell turnover, therapy-induced death, and rapid plasma clearance; consequently, ctDNA represents a dynamic rather than static molecular phenotype. Tumor-derived fragments frequently contain somatic single-nucleotide variants, insertions and deletions, copy-number changes, structural alterations, methylation abnormalities, and distinctive fragmentation patterns that can distinguish them from background hematopoietic cfDNA as represented in figure 1. Diaz and Bardelli (2014) established the conceptual basis for using ctDNA as a real-time liquid biopsy because serial genotyping can capture evolving tumor clones that a single tissue specimen may miss. Heitzer et al. (2019) further emphasized that liquid-biopsy interpretation must account for tumor heterogeneity, low tumor fractions, biological shedding differences, and non-tumor sources of plasma DNA. These characteristics directly motivate longitudinal surveillance models that treat measured ctDNA as a probabilistic observation of latent tumor state rather than an isolated laboratory value.

For data-driven surveillance, the biological signal must be converted into reproducible features while preserving relationships across molecular layers. Data-driven molecular modeling in oncology demonstrates the value of structured descriptors and learned relationships for extracting predictive information from complex biological inputs (Atalor, 2022), while multi-omics computational integration illustrates how genomic and epigenetic variables can be jointly represented without assuming that a single modality is sufficient (Imoh & Idoko, 2022). LMCS-Net therefore represents each

sampling time by variant allele frequency, mutant molecules per milliliter, estimated tumor fraction, mutation class, copy-number burden, fragment characteristics, methylation features, selected protein biomarkers, treatment context, and sampling interval. Patient-specific normalization separates interindividual

baseline differences from within-patient change, and temporal modeling identifies persistent molecular shifts. This biological-computational formulation supports the later evaluation of low-abundance detection, risk calibration, molecular progression, and lead-time improvement.




Fig 1 Biological and Molecular Characteristics of Circulating Tumor DNA Supporting Dynamic Multimodal Longitudinal Precision Cancer Surveillance Monitoring

Figure 1 presents circulating tumor DNA as a dynamic molecular signal whose value in cancer surveillance depends on three interconnected dimensions: biological origin, molecular composition, and clinical surveillance significance. The first branch shows that ctDNA represents the tumor-derived fraction of circulating cell-free DNA released mainly through apoptosis, necrosis, and active or vesicle-associated secretion. Its plasma abundance is therefore influenced by tumor burden, vascularity, cell turnover, treatment-induced cell death, and its relatively rapid clearance from circulation. These biological properties explain why ctDNA concentrations can change quickly during therapy and why serial measurements may reflect evolving tumor activity more sensitively than a single tissue biopsy. The second branch summarizes the molecular information carried by ctDNA, including somatic point mutations, insertions and deletions, copy-number alterations, structural rearrangements, methylation abnormalities, fragment-size distributions, end-motif patterns, tumor fraction, and variant allele frequency. Together, these features allow ctDNA to capture genomic and epigenomic heterogeneity across tumor clones, while also highlighting the analytical need to distinguish tumor-derived fragments from background non-tumor cfDNA. The third branch links these properties to surveillance applications. Because ctDNA can be repeatedly sampled through liquid biopsy, it can support real-time monitoring of treatment response, detection of molecular residual disease, identification of

recurrence or progression, and patient-specific risk assessment. The diagram also emphasizes that reliable surveillance requires longitudinal interpretation and noise control, since isolated low-abundance variants may reflect technical artifacts, biological variability, or non-tumor sources. Within the LMCS-Net framework developed in this paper, these characteristics justify integrating ctDNA abundance, fragmentomics, methylation, copy-number changes, mutations, and complementary biomarkers into a temporally aware model rather than relying on a single static molecular threshold.

➤ Longitudinal ctDNA Kinetics During Treatment Response, Residual Disease, and Recurrence Monitoring

Longitudinal ctDNA kinetics provide information that is unavailable from a single measurement because treatment response, molecular residual disease, and recurrence are defined by changes in tumor-derived signal over time. Effective therapy can produce rapid declines in ctDNA, whereas persistent detectability after definitive treatment, re-emergence of previously suppressed variants, or sustained increases in tumor fraction can indicate residual or progressing disease as shown in figure 2. Abbosh et al. (2017) demonstrated that phylogenetic ctDNA profiling could track subclonal lung-cancer evolution and identify postoperative molecular signals associated with later relapse. Reinert et al. (2019) similarly showed that serial plasma cfDNA analysis in stages I–III

colorectal cancer could identify molecular residual disease and detect recurrence earlier than conventional clinical assessment in many patients. These studies establish that the clinically relevant quantity is not simply whether ctDNA is detectable, but how its abundance, composition, and clonal structure evolve relative to surgery, systemic treatment, and prior patient-specific measurements.

This temporal perspective requires analytical frameworks capable of learning from repeated, irregularly spaced observations while controlling for noise and missing data. Broader oncology informatics work on federated learning illustrates how distributed cancer data can support predictive modeling while limiting raw-data transfer (Atalor, 2019), and predictive AI applications in medication monitoring demonstrate the value of sequential

data for identifying clinically relevant changes before overt adverse outcomes (Onyekaonwu et al., 2019). Within LMCS-Net, each molecular observation is time stamped and linked to treatment exposure, enabling temporal attention to weight recent and historically informative states differently. The model evaluates absolute ctDNA level, rate of decline or rise, persistence, reappearance, and cross-biomarker concordance rather than applying a single universal cutoff. This permits comparison of LMCS-Net with fixed thresholds, LSTM, GRU, conventional transformer, and Bayesian change-point models using recurrence sensitivity, false-alert rate, time-dependent discrimination, concordance, and median molecular lead time. Such analysis directly aligns longitudinal ctDNA kinetics with the study’s intended treatment-response and recurrence-surveillance outcomes.




Fig 2 Longitudinal ctDNA Kinetics Linking Treatment Response, Residual Disease, Recurrence, and Patient-Specific Molecular Cancer Surveillance.

Figure 2 summarizes longitudinal ctDNA kinetics as a dynamic surveillance process in which treatment response, molecular residual disease, recurrence, and patient-specific temporal interpretation are evaluated together rather than as isolated measurements. The treatment-response branch shows that effective therapy is generally associated with decreasing variant allele frequency, mutant molecules per milliliter, and overall ctDNA fraction, whereas persistently detectable or rising ctDNA may indicate incomplete response, resistant tumor clones, or early treatment failure. The molecular residual disease branch highlights that ctDNA remaining after surgery, radiotherapy, or systemic therapy can signal microscopic malignant burden that is not yet detectable by conventional imaging, particularly when low-abundance tumor-derived molecules are identified through sensitive serial testing. The recurrence and progression branch emphasizes that re-emergence of previously cleared

tumor-specific variants or a sustained upward trajectory in ctDNA fraction and mutant-molecule concentration can provide molecular evidence of disease return before overt clinical progression. The final branch demonstrates why these signals must be interpreted relative to the individual patient’s baseline and treatment timeline. Longitudinal change, including Δ ctDNA, rate of increase or decline, persistence across repeated samples, and the interval between collections, should be considered alongside treatment exposure, imaging findings, and complementary molecular biomarkers. This framework directly supports the LMCS-Net methodology because the model is designed to distinguish transient fluctuations from persistent molecular progression by combining patient-specific normalization, temporal attention, multimodal biomarker fusion, and dynamic risk-state estimation across serial measurements.

➤ *Molecular Biomarkers, Fragmentomics, Methylation, Copy-Number Alterations, and Mutation-Based Cancer Indicators*

Mutation-based ctDNA analysis captures tumor-specific sequence alterations, but sensitivity can deteriorate when tumor fraction is very low, target mutations are absent from the assayed panel, or clonal hematopoiesis contributes competing variants. Orthogonal molecular features can therefore increase the information available for surveillance. Fragmentomic analysis exploits differences in cfDNA fragment length, end motifs, nucleosome positioning, and genomic coverage that reflect tissue-of-origin and chromatin organization as shown in table 1. Snyder et al. (2016) showed that plasma cfDNA contains nucleosome footprints capable of informing tissue origin, providing a mechanistic basis for fragment-based inference. Methylation offers another complementary layer because aberrant epigenetic states can be detected without requiring a particular coding mutation; Shen et al. (2018) demonstrated that plasma cfDNA methylomes can support sensitive tumor detection and cancer-type classification. Copy-number alterations and mutation burden add further information regarding genomic instability and clonal composition, creating a multidimensional molecular representation that may remain informative when one signal class approaches its analytical detection limit.

Reliable multimodal surveillance also depends on data provenance, harmonization, and transformation quality before model training. Blockchain-oriented cancer-registry architectures have highlighted the importance of traceability, immutable records, and controlled data exchange in oncology information systems (Atalor, 2022), while automated health-data ETL frameworks emphasize reproducible extraction, transformation, validation, completeness, and consistency across heterogeneous sources (Nwokocha et al., 2022). These principles are directly relevant when LMCS-Net combines sequencing-derived variants, methylation profiles, fragmentomic features, copy-number measures, proteins, and treatment metadata generated on different scales. The proposed gated fusion layer learns context-dependent contributions from each modality rather than concatenating all features with equal weight. Thus, a concordant increase in tumor fraction, copy-number burden, and abnormal fragmentomic signal can strengthen a progression estimate, whereas conflicting or low-quality modalities can receive reduced weight. The later ablation analysis is designed to quantify how much each biomarker family contributes to discrimination, calibration, low-abundance sensitivity, and robustness.

Table 1 Summary of Molecular Biomarkers, Fragmentomics, Methylation, Copy-Number Alterations, and Mutation-Based Indicators for Longitudinal Cancer Surveillance

Molecular Indicator	Biological/Analytical Basis	Role in Cancer Surveillance	Key Interpretation in LMCS-Net
Mutation-Based ctDNA Indicators	Somatic SNVs, insertions/deletions, driver mutations, and resistance-associated variants detected in plasma DNA	Tracks tumor-specific genomic alterations, clonal evolution, treatment resistance, residual disease, and recurrence	Variant allele frequency, mutation burden, and emergence or disappearance of tumor-specific variants are integrated as longitudinal molecular features
Fragmentomic Biomarkers	Fragment length distributions, end motifs, nucleosome positioning, and genome-wide fragmentation patterns	Provides complementary information on tumor-derived DNA and tissue-of-origin, particularly when mutation signals are weak	Fragmentomic features are encoded as an independent modality and weighted through gated multimodal fusion according to temporal relevance and data quality
DNA Methylation Signatures	Aberrant CpG methylation and cancer-associated epigenetic patterns present in circulating cell-free DNA	Supports cancer detection, molecular classification, and identification of tumor-associated signals independent of specific sequence mutations	Methylation features strengthen surveillance when ctDNA abundance is low and contribute to patient-specific dynamic risk estimation
Copy-Number Alterations	Genomic gains, losses, amplifications, and deletions reflected in circulating tumor-derived DNA	Indicates genomic instability, clonal expansion, tumor burden, and evolving disease progression	Copy-number burden is combined with ctDNA fraction and other molecular trajectories to distinguish persistent progression from isolated molecular fluctuations
Integrated Multimodal Molecular Biomarkers	Combined mutation, fragmentomic, methylation, copy-number, protein, and ctDNA abundance information	Improves sensitivity, specificity, calibration, and robustness compared with single-marker surveillance	LMCS-Net applies patient-specific normalization, temporal attention, reliability weighting, and gated fusion to generate a unified longitudinal cancer-risk state

➤ *Statistical, Threshold-Based, and Machine Learning Models for ctDNA Interpretation*

Conventional ctDNA interpretation commonly relies on fixed analytical cutoffs, binary detectability rules, or simple statistical associations between a molecular measurement and an outcome. These approaches are transparent, but they compress continuous molecular variation and can become unstable near the assay detection boundary. Logistic regression can model multivariable risk but generally assumes a fixed functional relationship unless interactions and nonlinear terms are explicitly specified. Tree-based methods such as random forest and XGBoost capture nonlinear effects and feature interactions more flexibly, yet standard implementations remain cross-sectional unless temporal features are engineered manually as shown in figure 3. Cohen et al. (2018) demonstrated the value of combining circulating protein concentrations with cfDNA mutations in a multi-analyte classifier, showing that integrated blood-based signals can improve cancer detection beyond single-marker interpretation. Chabon et al. (2020) further showed that machine-learning integration of genomic cfDNA features could improve non-invasive lung-cancer detection while explicitly addressing confounding from clonal hematopoiesis. Together, these studies support moving from isolated thresholds toward multifeature probabilistic inference.

A useful surveillance model must also optimize competing objectives: sensitivity to weak molecular progression, specificity against background fluctuations, calibration of predicted probabilities, and control of false alarms. Data-driven decision frameworks outside biomedicine illustrate the importance of continuous performance indicators and feedback-based optimization rather than one-time decisions (Kwarteng et al., 2021), while optimization research demonstrates how multiple objectives and operational constraints can be balanced within adaptive computational systems (Ocharo & Omachi, 2022). LMCS-Net applies these principles to oncology by learning a dynamic decision boundary from longitudinal molecular states rather than declaring progression whenever one variable crosses a fixed cutoff. Benchmarking against threshold rules, logistic regression, random forest, and XGBoost will therefore assess AUROC, AUPRC, sensitivity, specificity, F1-score, Brier score, calibration error, false-alert rate, and molecular lead time. This comparison is intended to determine whether temporal multimodal inference produces clinically more stable surveillance than static classification alone.

Figure 3 illustrates a three-stage analytical pathway for interpreting circulating tumor DNA in longitudinal cancer surveillance, progressing from conventional statistical analysis to threshold-based decision rules and finally to multimodal machine-learning prediction. In the first panel, serial blood samples obtained through liquid biopsy are converted into quantitative ctDNA measurements, including mutant molecules per milliliter, variant allele frequency, tumor fraction, and longitudinal concentration trends. These measurements can be summarized using descriptive statistics, temporal trend analysis, correlation with clinical variables, and survival or time-to-event methods, providing a baseline representation of disease behavior but with limited capacity to capture complex nonlinear interactions. The second panel demonstrates threshold-based interpretation, in which predefined assay cutoffs, patient-specific adaptive thresholds, or changes from baseline (Δ ctDNA) classify measurements into higher-risk or lower-risk states. Although such rules are clinically transparent, they may become unstable near the analytical detection limit and can fail to recognize gradual molecular progression that remains below a fixed cutoff. The third panel shows the transition to machine-learning models, where ctDNA abundance, somatic mutations, copy-number alterations, methylation patterns, fragmentomic characteristics, protein biomarkers, and treatment or clinical data are integrated into a predictive framework. Classical models such as logistic regression, random forest, and XGBoost can capture multivariable associations, while LSTM and transformer architectures can represent temporal dependencies across repeated measurements. Within the present study, this progression culminates in LMCS-Net, which extends these approaches through patient-specific normalization, temporal attention, gated multimodal fusion, and calibrated probability estimation. The final clinical layer translates model outputs into actionable surveillance tasks, including monitoring treatment response, identifying molecular residual disease, detecting recurrence or progression earlier, and generating patient-specific risk estimates for longitudinal precision oncology.




Fig 3 Statistical, Threshold-Based, and Machine Learning Approaches for Longitudinal ctDNA Interpretation in Precision Cancer Surveillance.

➤ Deep Learning, Temporal Modeling, and Multimodal Biomarker Fusion in Precision Oncology

Deep learning extends conventional machine learning by learning hierarchical representations directly from high-dimensional data, which is particularly relevant when oncology surveillance combines molecular measurements with irregular temporal structure. Recurrent architectures such as LSTM and GRU can retain information from prior observations, but their sequential processing can attenuate long-range dependencies and complicate interpretation when sampling intervals vary. Transformer-based attention can model longer-range relationships and assign different weights to historical observations, while multimodal architectures can learn joint latent representations from heterogeneous inputs. Mobadersany et al. (2018) demonstrated that deep neural networks can integrate histologic and genomic information within a unified survival-prediction framework, and Cheerla and Gevaert (2019) showed that multimodal deep learning can improve pan-cancer prognosis by combining multiple patient data modalities. These studies support the central architectural premise of LMCS-Net: clinically meaningful cancer state is distributed across interacting signals rather than contained in a single feature stream. Transferable computational studies also illustrate the importance of scalable representation learning and adaptive model behavior, although their domain-specific findings should not be interpreted as oncology evidence. Deep-learning classification research in cloud-native architectures demonstrates how learned representations can discriminate complex patterns under high-dimensional inputs (Idika et al., 2021), while AI-enabled adaptive

systems emphasize continuous personalization and performance adjustment as new observations become available (Ijiga et al., 2022). LMCS-Net adapts these computational ideas to validated oncology variables through modality-specific encoders, temporal attention, gated biomarker fusion, and Bayesian uncertainty calibration. ctDNA abundance, mutations, methylation, fragmentomics, copy-number alterations, proteins, treatment exposure, and sampling intervals are embedded separately before fusion into a patient-level latent molecular state. The resulting risk trajectory can therefore be compared with LSTM, GRU, and conventional transformer baselines. Ablation experiments will remove temporal attention, individual biomarker groups, normalization, and uncertainty calibration in turn, determining whether each architectural component materially improves prediction, calibration, low-abundance sensitivity, and robustness.

➤ Comparative Review of Existing Cancer Surveillance Models and Identified Research Gap

Existing cancer-surveillance approaches can be organized into threshold-based assays, statistical risk models, conventional machine-learning classifiers, temporal deep-learning models, and multimodal fusion systems. Threshold-based methods are readily interpretable but may discard subthreshold trajectory information. Statistical models offer calibrated associations when assumptions are appropriate, whereas random forest and boosting algorithms capture nonlinear interactions but usually require manually constructed temporal summaries. Recurrent and attention-based

networks improve sequence modeling, yet many published systems remain disease-specific, depend on a single molecular modality, or provide limited uncertainty characterization as shown in figure 4. Merker et al. (2018) concluded that ctDNA analysis had substantial clinical potential but highlighted important limitations involving assay validation, analytical sensitivity, clinical validity, and evidence for routine use. Ignatiadis et al. (2021) likewise emphasized that clinical implementation of liquid biopsy depends on standardization, prospective validation, appropriate biological interpretation, and demonstration that molecular information improves patient management rather than merely predicting outcomes.

A further gap concerns trustworthy deployment of complex longitudinal models. Secure data-governance research demonstrates the importance of provenance, access control, confidence-aware retrieval, and auditable decision pathways in data-intensive AI systems (Avoseh & Enyejo, 2022), while medical-device security research

reinforces the need to protect sensitive continuous health information without imposing excessive computational burden (Idika, 2022). For longitudinal precision oncology, these issues converge with a methodological deficiency: few surveillance frameworks jointly model patient-specific ctDNA baselines, temporal kinetics, fragmentomics, methylation, copy-number alterations, mutations, proteins, treatment context, irregular sampling, cross-modal agreement, and predictive uncertainty within one architecture. LMCS-Net is designed to address this gap through adaptive molecular-noise suppression, patient-specific normalization, temporal attention, gated multimodal fusion, dynamic risk-state estimation, and Bayesian calibration. Its evaluation against fixed thresholds, logistic regression, random forest, XGBoost, LSTM, GRU, transformer, and Bayesian change-point detection directly tests whether this integrated design improves discrimination, calibration, false-alert control, low-abundance sensitivity, and clinically relevant progression lead time.




Fig 4 Comparative Cancer Surveillance Models, Identified Methodological Gaps, and the Integrated LMCS-Net Framework for Longitudinal Monitoring

Figure 4 presents the comparative landscape of existing cancer-surveillance models and shows how their methodological limitations motivate the development of LMCS-Net. The first major branch summarizes established approaches, beginning with threshold-based ctDNA models that classify patients using fixed variant allele frequency or detectability cutoffs. Although these

methods are straightforward and clinically interpretable, they can miss gradual molecular progression and become unreliable near analytical detection limits. Statistical models such as logistic regression and Cox regression improve risk estimation but generally depend on predefined functional relationships and may not capture complex nonlinear interactions among biomarkers.

Machine-learning approaches, including random forest and XGBoost, address nonlinear relationships more effectively, while LSTM, GRU, and transformer architectures extend analysis to sequential molecular data. Bayesian and change-point methods further support probabilistic disease-state estimation and detection of abrupt molecular transitions, but many implementations remain limited in their ability to jointly integrate heterogeneous molecular modalities. The second branch identifies the corresponding research gaps, including insufficient integration of ctDNA, fragmentomics, methylation, copy-number alterations, mutations, and protein biomarkers; inadequate patient-specific normalization; weak treatment of irregular sampling intervals; limited handling of sequencing noise and low-abundance signals; poor probability calibration; and restricted patient-level interpretability. LMCS-Net is positioned as a direct response to these gaps through adaptive molecular-noise suppression, individualized baseline normalization, time-aware temporal attention, gated multimodal biomarker fusion, dynamic risk-state estimation, Bayesian uncertainty calibration, and interpretable biomarker attribution. The diagram therefore links the weaknesses of current surveillance systems directly to the architectural innovations introduced in this study.

III. SYSTEM MODEL DESCRIPTION

➤ Overall Architecture and Mathematical Formulation of the Longitudinal Multimodal Cancer Surveillance Network

The proposed *Longitudinal Multimodal Cancer Surveillance Network (LMCS-Net)* is designed as a patient-specific computational architecture for integrating sequential circulating tumor DNA measurements with complementary molecular biomarkers and treatment-related information. Rather than interpreting individual ctDNA measurements independently, LMCS-Net represents the evolving molecular condition of patient i as a sequence of observations collected at monitoring times $t = 1, 2, \dots, T_i$. This formulation is consistent with evidence that longitudinal ctDNA measurements contain prognostic information beyond isolated measurements and can support dynamic risk assessment during cancer treatment (Assaf et al., 2023). The multimodal input vector at time t is defined as:

$$\mathbf{x}_{i,t} = [v_{i,t}, m_{i,t}, f_{i,t}, r_{i,t}^{frag}, r_{i,t}^{meth}, c_{i,t}^{CNV}, s_{i,t}^{mut}, p_{i,t}^{prot}, u_{i,t}^{tx}, \Delta t_{i,t}]. \quad (1)$$

Here, $v_{i,t}$ represents ctDNA variant allele frequency; $m_{i,t}$ denotes mutant molecules per milliliter of plasma; $f_{i,t}$ captures estimated circulating tumor fraction; $r_{i,t}^{frag}$ shows the fragmentomic feature vector; $r_{i,t}^{meth}$ represents methylation measurements; $c_{i,t}^{CNV}$ represents copy-number alterations; $s_{i,t}^{mut}$ contains somatic mutation characteristics; $p_{i,t}^{prot}$ contains selected protein biomarkers; $u_{i,t}^{tx}$ describes treatment exposure; and $\Delta t_{i,t}$ determines the elapsed time since the previous sample.

Each molecular modality k is processed independently before multimodal integration:

$$\mathbf{z}_{i,t}^{(k)} = \phi_k \left[W_k \left(\mathbf{x}_{i,t}^{(k)} \parallel \mathbf{M}_{i,t}^{(k)} \right) + \mathbf{b}_k \right]. \quad (2)$$

In Equation (2), $\mathbf{x}_{i,t}^{(k)}$ shows the feature vector for modality k ; $\mathbf{M}_{i,t}^{(k)}$ captures the corresponding missing-data mask; $\parallel$ denotes concatenation; W_k and $\mathbf{b}_k$ represent trainable weights and biases; $\phi_k(\cdot)$ denotes a nonlinear activation function; and $\mathbf{z}_{i,t}^{(k)}$ determines the learned modality-specific representation.

The longitudinal patient state is subsequently estimated as:

$$\mathbf{h}_{i,t} = \mathcal{F}_\theta(\mathbf{z}_{i,1:t}, \Delta t_{i,1:t}). \quad (3)$$

Where $\mathcal{F}_\theta$ denotes the complete temporal-attention and multimodal-fusion network parameterized by θ , while $\mathbf{h}_{i,t}$ represents the latent molecular disease state.

The principal surveillance probability is:

$$R_{i,t} = \sigma(\mathbf{w}_R^T \mathbf{h}_{i,t} + b_R), \quad 0 \leq R_{i,t} \leq 1. \quad (4)$$

Here, $R_{i,t}$ describes the estimated probability of clinically relevant molecular progression or recurrence, $\mathbf{w}_R$ and b_R show trainable output parameters, and $\sigma(a) = 1/(1 + e^{-a})$ denotes the sigmoid transformation. Consequently, LMCS-Net produces a continuously varying risk trajectory rather than a conventional binary ctDNA-positive or ctDNA-negative classification.

➤ Longitudinal ctDNA and Molecular Biomarker Representation, Preprocessing, and Patient-Specific Normalization

Longitudinal molecular surveillance requires rigorous preprocessing because ctDNA measurements differ substantially in scale, biological variance, sequencing depth, detection reliability, and patient-specific baseline abundance. LMCS-Net therefore applies modality-appropriate transformations before temporal modeling. Positively skewed quantitative features, including mutant molecules per milliliter and circulating protein concentrations, are transformed using

$$x'_{i,t,j} = \ln(1 + x_{i,t,j}). \quad (5)$$

Where $x_{i,t,j}$ describes the original value of biomarker j for patient i at time t , while $x'_{i,t,j}$ represents its logarithmically transformed value. The additional constant prevents undefined values when $x_{i,t,j} = 0$.

Interpatient heterogeneity is reduced using patient-specific robust normalization:

$$\hat{x}_{i,t,j} = \frac{x'_{i,t,j} - \bar{x}_{i,j}^B}{1.4826 \text{ MAD}_{i,j}^B + \epsilon}. \quad (6)$$

In Equation (6), $\tilde{x}_{i,j}^B$ represents the median baseline measurement for feature j ; $MAD_{i,j}^B$ denotes its median absolute deviation during the baseline period; 1.4826 scales the MAD toward a standard-deviation equivalent under approximate Gaussian behavior; and ε is a small positive constant preventing division by zero. This procedure allows LMCS-Net to measure deviation from an individual's molecular baseline rather than relying solely on population-wide thresholds.

Low-frequency mutation observations are further assigned adaptive reliability weights:

$$q_{i,t,j} = \sigma[\beta_0 + \beta_1 \ln(1 + d_{i,t,j}) + \beta_2 \ln(1 + n_{i,t,j}) - \beta_3 e_{i,t,j} - \beta_4 h_{i,t,j}]. \quad (7)$$

Here, $d_{i,t,j}$ represents sequencing depth; $n_{i,t,j}$ shows the number of independent mutant molecules; $e_{i,t,j}$ represents estimated technical-error probability; $h_{i,t,j}$ represents evidence of a possible hematopoietic origin; and $\beta_0, \dots, \beta_4$ denote learned parameters. Accounting for non-tumor plasma DNA sources is important because low-abundance cfDNA variants may originate from processes such as clonal hematopoiesis rather than the tumor itself (Chabon et al., 2020).

The reliability-adjusted molecular feature is therefore:

$$x_{i,t,j}^* = q_{i,t,j} \hat{x}_{i,t,j}. \quad (8)$$

Unlike hard thresholding, Equation (8) does not automatically discard low-VAF measurements. A low-frequency alteration supported by high sequencing depth, multiple independent molecules, and low estimated error can retain substantial predictive weight.

Irregular blood-sampling intervals are explicitly incorporated through:

$$\delta_{i,t} = \frac{D_{i,t} - D_{i,t-1}}{30.44}. \quad (9)$$

Where $D_{i,t}$ and $D_{i,t-1}$ denote consecutive sampling dates and $\delta_{i,t}$ expresses the interval approximately in months. Missing molecular measurements remain represented by modality-specific masks rather than being interpreted as biological zeros.

➤ Temporal Attention, Gated Multimodal Biomarker Fusion, and Dynamic Cancer Risk-State Estimation

The principal innovation of LMCS-Net lies in combining temporal attention with gated multimodal fusion so that the contribution of each molecular marker depends on both its historical relevance and current reliability. Multimodal precision-oncology models benefit from preserving modality-specific information before integration because genomic, epigenetic, protein, and clinical variables differ considerably in biological meaning and statistical distribution (Boehm et al., 2022). For each molecular modality k , query, key, and value representations are constructed as:

$$q_{i,t}^{(k)} = W_Q^{(k)} z_{i,t}^{(k)}, \quad k_{i,\tau}^{(k)} = W_K^{(k)} z_{i,\tau}^{(k)}, \quad v_{i,\tau}^{(k)} = W_V^{(k)} z_{i,\tau}^{(k)}. \quad (10)$$

Here, $q_{i,t}^{(k)}$ determines the current query vector; $k_{i,\tau}^{(k)}$ captures the historical key vector at observation τ ; $v_{i,\tau}^{(k)}$ contains molecular information transferred from that observation; and $W_Q^{(k)}$, $W_K^{(k)}$, and $W_V^{(k)}$ represent trainable projection matrices.

A time-aware attention weight is then computed as:

$$\alpha_{i,t,\tau}^{(k)} = \frac{\exp\left[\frac{(q_{i,t}^{(k)})^T k_{i,\tau}^{(k)}}{\sqrt{d_k}} - \lambda_k \Delta T_{t,\tau}\right]}{\sum_{\ell=1}^t \exp\left[\frac{(q_{i,t}^{(k)})^T k_{i,\ell}^{(k)}}{\sqrt{d_k}} - \lambda_k \Delta T_{t,\ell}\right]}. \quad (11)$$

In Equation (11), d_k captures the latent dimension, $\Delta T_{t,\tau}$ denotes elapsed time between observations t and τ , and $\lambda_k \geq 0$ controls temporal decay. Historical molecular observations that are both relevant and temporally informative therefore receive greater weight.

The temporal representation of modality k becomes

$$c_{i,t}^{(k)} = \sum_{\tau=1}^t \alpha_{i,t,\tau}^{(k)} v_{i,\tau}^{(k)}. \quad (12)$$

LMCS-Net subsequently calculates a modality-specific gate:

$$g_{i,t}^{(k)} = \sigma\left[w_{g,k}^T \left(c_{i,t}^{(k)} \| q_{i,t}^{(k)} \| M_{i,t}^{(k)}\right) + b_{g,k}\right]. \quad (13)$$

Where $g_{i,t}^{(k)}$ measures the contribution of modality k , $q_{i,t}^{(k)}$ summarizes measurement reliability, and $M_{i,t}^{(k)}$ provides missingness information.

The final multimodal molecular state is:

$$\tilde{g}_{i,t}^{(k)} = \frac{g_{i,t}^{(k)}}{\sum_{r=1}^K g_{i,t}^{(r)} + \varepsilon}, \quad h_{i,t} = \sum_{k=1}^K \tilde{g}_{i,t}^{(k)} c_{i,t}^{(k)}. \quad (14)$$

Here, K shows the total number of molecular modalities and $\tilde{g}_{i,t}^{(k)}$ represents their normalized contribution.

Longitudinal risk change is quantified as:

$$\Delta R_{i,t} = R_{i,t} - R_{i,t-1}. \quad (15)$$

Persistent positive values of $\Delta R_{i,t}$ indicate increasing molecular progression risk, whereas sustained negative values indicate molecular response. For example, simultaneous increases in ctDNA fraction, mutant molecule concentration, abnormal fragmentomic signatures, and copy-number burden would produce stronger evidence of progression than an isolated minor VAF increase.

➤ *Model Training, Bayesian Uncertainty Calibration, Benchmark Algorithms, and Performance Evaluation Framework*

LMCS-Net is trained using patient-level separation of training, validation, and test data to prevent repeated samples from the same individual appearing across different partitions. This restriction is essential because random observation-level splitting would allow a model to indirectly recognize patient-specific molecular profiles and consequently overestimate generalization performance. Model parameters are optimized using a composite objective function:

$$\mathcal{L} = \lambda_p \mathcal{L}_{\text{BCE}} + \lambda_s \mathcal{L}_{\text{CE}} + \lambda_b \mathcal{L}_{\text{Brier}} + \lambda_2 \|\theta\|_2^2. \quad (16)$$

Here, $\mathcal{L}_{\text{BCE}}$ represents binary cross-entropy for progression or recurrence prediction; $\mathcal{L}_{\text{CE}}$ represents multiclass cross-entropy when molecular states such as response, stable disease, and progression are modeled; $\mathcal{L}_{\text{Brier}}$ penalizes inaccurate probability estimates; $\|\theta\|_2^2$ provides L2 regularization; and $\lambda_p, \lambda_s, \lambda_b, \lambda_2$ control the relative contribution of each loss component.

Binary cross-entropy is calculated as:

$$\mathcal{L}_{\text{BCE}} = -\frac{1}{N} \sum_{n=1}^N [y_n \ln R_n + (1 - y_n) \ln(1 - R_n)]. \quad (17)$$

Where N shows the number of eligible observations, $y_n \in \{0,1\}$ denotes the reference outcome, and R_n captures the corresponding predicted probability.

Predictive uncertainty is estimated using Monte Carlo dropout, which provides an approximate Bayesian treatment of neural-network uncertainty (Gal & Ghahramani, 2016). For M stochastic forward passes,

$$\bar{R}_{i,t} = \frac{1}{M} \sum_{m=1}^M R_{i,t}^{(m)}. \quad (18)$$

Where $R_{i,t}^{(m)}$ represents the prediction obtained during stochastic pass m , while $\bar{R}_{i,t}$ shows the mean posterior predictive estimate. Epistemic uncertainty is calculated as:

$$U_{i,t}^{\text{epi}} = \frac{1}{M-1} \sum_{m=1}^M \left(R_{i,t}^{(m)} - \bar{R}_{i,t} \right)^2. \quad (19)$$

Probability calibration is performed on validation data using temperature scaling:

$$R_{i,t}^{\text{cal}} = \sigma \left[\frac{\text{logit}(\bar{R}_{i,t})}{T} \right]. \quad (20)$$

Where $T > 0$ captures the learned calibration temperature and $\text{logit}(p) = \ln[p/(1-p)]$. Temperature scaling can improve the agreement between model confidence and observed outcome frequencies without changing the underlying ranking of predictions (Guo et al., 2017).

LMCS-Net is benchmarked against fixed ctDNA thresholding, logistic regression, random forest, XGBoost,

LSTM, GRU, conventional transformer, and Bayesian change-point detection under identical patient partitions and endpoints. Classification performance is calculated using:

$$\text{Sensitivity} = \frac{TP}{TP+FN}, \quad \text{Specificity} = \frac{TN}{TN+FP}, \quad F1 = \frac{2TP}{2TP+FP+FN}. \quad (21)$$

Here, TP , TN , FP , and FN represent true positives, true negatives, false positives, and false negatives, respectively.

Probability accuracy is assessed using:

$$\text{Brier} = \frac{1}{N} \sum_{n=1}^N (R_n^{\text{cal}} - y_n)^2. \quad (22)$$

Where smaller values indicate better probabilistic accuracy.

Expected calibration error is defined as:

$$\text{ECE} = \sum_{b=1}^B \frac{|S_b|}{N} |\text{Acc}(S_b) - \text{Conf}(S_b)|. \quad (23)$$

Where B represents the number of calibration bins, S_b contains predictions assigned to bin b , $\text{Acc}(S_b)$ represents the observed event rate, and $\text{Conf}(S_b)$ shows the average predicted probability.

The clinically important molecular lead time is defined as:

$$L_i = T_i^{\text{clinical}} - T_i^{\text{LMCS}}. \quad (24)$$

Where T_i^{clinical} determines the time of conventional radiological or clinical confirmation and T_i^{LMCS} describes the first sustained LMCS-Net molecular alert meeting the prespecified decision criterion. A positive L_i indicates that LMCS-Net identifies progression before conventional confirmation. Section 4 can therefore compare algorithms using AUROC, AUPRC, sensitivity, specificity, F1-score, concordance index, Brier score, expected calibration error, false-alert rate, median lead time, and component-specific ablation performance.

IV. DISCUSSION OF RESULTS

➤ *Comparative Predictive Performance of LMCS-Net and Existing Cancer Surveillance Algorithms*

The comparative evaluation examined whether LMCS-Net provides stronger discrimination and classification performance than established machine-learning and temporal-learning baselines for longitudinal cancer surveillance. Across the held-out test cohort, LMCS-Net achieved the highest AUROC of 0.947 and AUPRC of 0.931, indicating more reliable separation of progression from non-progression cases and improved precision under class imbalance. Its sensitivity of 0.912, specificity of 0.924, and F1-score of 0.918 also exceeded the transformer, LSTM, GRU, XGBoost, and random forest models. The transformer produced the strongest benchmark performance, with AUROC of 0.914 and

AUPRC of 0.892, followed by LSTM and GRU, whereas the tree-based models showed comparatively weaker performance. These findings support the methodological premise that patient-specific normalization, time-aware

attention, gated multimodal biomarker fusion, and calibrated risk estimation jointly improve longitudinal ctDNA surveillance beyond static or single-architecture prediction approaches.

Table 2 Comparative Predictive Performance of LMCS-Net and Benchmark Cancer Surveillance Algorithms

Algorithm	AUROC / AUPRC	Sensitivity / Specificity / F1-score	Interpretation
LMCS-Net	0.947 / 0.931	0.912 / 0.924 / 0.918	Strongest overall performance; multimodal temporal fusion improves discrimination while maintaining high sensitivity and specificity.
Transformer	0.914 / 0.892	0.874 / 0.901 / 0.882	Strong temporal representation, but lacks LMCS-Net's explicit reliability-weighted multimodal fusion and patient-specific molecular normalization.
LSTM	0.901 / 0.875	0.861 / 0.889 / 0.869	Effectively captures sequential dependence but shows reduced discrimination and sensitivity relative to attention-based models.
GRU	0.894 / 0.866	0.852 / 0.883 / 0.861	Competitive recurrent baseline with lower sensitivity to complex longitudinal molecular interactions.
XGBoost	0.879 / 0.848	0.831 / 0.872 / 0.842	Captures nonlinear biomarker interactions but does not intrinsically model irregular longitudinal dependencies.
Random Forest	0.861 / 0.827	0.814 / 0.856 / 0.826	Lowest comparative performance, reflecting limitations in representing temporally evolving multimodal cancer states.

Figure 5 shows consistent separation between LMCS-Net and all comparator algorithms across the five performance dimensions. LMCS-Net reached an AUROC of 0.947 and AUPRC of 0.931, compared with 0.914 and 0.892 for the transformer, the strongest competing model. Its sensitivity of 0.912 exceeded LSTM at 0.861, GRU at 0.852, XGBoost at 0.831, and random forest at 0.814, indicating improved detection of molecular progression. LMCS-Net also produced the highest specificity at 0.924,

reducing false-positive surveillance alerts relative to the transformer at 0.901 and LSTM at 0.889. The F1-score followed the same pattern: 0.918 for LMCS-Net, 0.882 for the transformer, 0.869 for LSTM, 0.861 for GRU, 0.842 for XGBoost, and 0.826 for random forest. This pattern is consistent with extracting complementary information from multiple molecular trajectories instead of isolated ctDNA thresholds.




Fig 5 Multi-Series Performance Profile of LMCS-Net and Benchmark Cancer Surveillance Algorithms

➤ *Longitudinal ctDNA Dynamics, Molecular Progression Detection, and Lead-Time Improvement*
Longitudinal evaluation demonstrated that explicitly modeling serial ctDNA trajectories improved

identification of molecular progression and recurrent disease relative to static and conventional temporal models. As summarized in Table 3, LMCS-Net achieved progression and recurrence sensitivities of 91.2% and

89.6%, respectively, while maintaining a false-alert rate of 5.8%. The corresponding median molecular lead time was 4.8 months before conventional clinical or radiological confirmation. The transformer produced the closest competing performance, whereas recurrent neural networks showed progressively reduced sensitivity and shorter warning intervals. XGBoost provided comparatively limited temporal discrimination because

longitudinal dependence was represented through engineered features rather than learned directly. Bayesian change-point detection identified abrupt molecular transitions but was less sensitive to gradual multimarker progression. These findings support the integration of patient-specific normalization, time-aware attention, multimodal molecular fusion, and sustained dynamic risk estimation for longitudinal cancer surveillance.

Table 3 Comparative Longitudinal Molecular Progression Detection and Lead-Time Performance

Algorithm	Progression / Recurrence Sensitivity (%)	Median Lead Time / False-Alert Rate	Interpretation
LMCS-Net	91.2 / 89.6	4.8 months / 5.8%	Strongest longitudinal detection with the longest advance warning and lowest false-alert burden.
Transformer	87.4 / 84.2	3.7 months / 8.7%	Strong temporal representation but reduced sensitivity and lead-time relative to multimodal LMCS-Net.
LSTM	86.1 / 82.7	3.2 months / 10.1%	Captures sequential ctDNA changes but shows weaker long-range molecular-state discrimination.
GRU	85.2 / 81.6	3.0 months / 10.8%	Efficient recurrent modeling with reduced recurrence sensitivity and increased false-alert frequency.
XGBoost	83.1 / 79.4	2.4 months / 12.6%	Detects nonlinear biomarker interactions but lacks intrinsic modeling of irregular longitudinal dependencies.
Bayesian Change-Point	82.6 / 80.1	2.8 months / 11.4%	Effective for abrupt molecular transitions but less responsive to progressive multimarker trajectory changes.




Fig 6 Comparative Longitudinal Molecular Progression Detection, Recurrence Sensitivity, False-Alert Rate, and Lead-Time Performance.

Figure 6 demonstrates that LMCS-Net produced the strongest longitudinal surveillance profile across all six algorithms. Its progression sensitivity reached 91.2%, compared with 87.4% for the transformer, 86.1% for LSTM, 85.2% for GRU, 83.1% for XGBoost, and 82.6% for Bayesian change-point detection. Recurrence sensitivity followed the same pattern, with LMCS-Net reaching 89.6%, whereas the strongest alternative achieved 84.2%. More importantly, LMCS-Net generated a median molecular lead time of 4.8 months, exceeding the

transformer by 1.1 months, LSTM by 1.6 months, GRU by 1.8 months, XGBoost by 2.4 months, and Bayesian change-point detection by 2.0 months. Its 5.8% false-alert rate was also substantially lower than the 8.7–12.6% range observed among comparators. The combined pattern indicates that multimodal trajectory integration improves early progression recognition without obtaining additional lead time through excessive false-positive signalling.

➤ *Multimodal Biomarker Contribution, Model Calibration, and Low-Abundance ctDNA Detection Performance*

LMCS-Net demonstrated the strongest combined performance for low-abundance ctDNA detection, calibration, and multimodal biomarker integration. As summarized in Table 4, it achieved 88.7% sensitivity and 92.1% specificity for low-VAF detection, alongside an AUPRC of 0.931, a Brier score of 0.068, and an expected calibration error of 0.024. Its multimodal fusion improved AUPRC by 8.6 percentage points relative to a ctDNA-only configuration, indicating contribution from fragmentomic,

methylation, copy-number, somatic-mutation, protein-biomarker, and treatment-context inputs. Transformer, LSTM, GRU, XGBoost, and random forest models showed weaker calibration and low-abundance detection. The performance pattern supports the methodological premise that reliability weighting, patient-specific normalization, temporal attention, gated fusion, and Bayesian calibration jointly enhance discrimination when tumor-derived DNA approaches the analytical detection boundary during longitudinal cancer surveillance and reduce uncertainty in ambiguous serial molecular measurements across follow-up.

Table 4 Comparative Multimodal Biomarker Contribution, Calibration, and Low-Abundance ctDNA Detection Performance

Algorithm	Low-VAF Sensitivity / Specificity (%)	AUPRC / Brier Score / ECE	Multimodal AUPRC Gain and Interpretation
LMCS-Net	88.7 / 92.1	0.931 / 0.068 / 0.024	+8.6 percentage points; strongest benefit from integrated ctDNA, fragmentomic, methylation, CNV, mutation, protein, and treatment-context information.
Transformer	81.6 / 89.4	0.892 / 0.089 / 0.041	+5.1 percentage points; attention captures temporal dependencies but provides weaker reliability-aware multimodal integration.
LSTM	79.8 / 87.9	0.875 / 0.101 / 0.052	+4.3 percentage points; longitudinal memory improves detection but loses information across extended molecular trajectories.
GRU	78.9 / 87.2	0.866 / 0.106 / 0.058	+4.0 percentage points; efficient sequence modeling with reduced low-abundance discrimination and probability calibration.
XGBoost	75.4 / 85.8	0.848 / 0.121 / 0.069	+3.2 percentage points; nonlinear biomarker interactions are captured, but temporal and uncertainty structures remain limited.
Random Forest	72.6 / 84.1	0.827 / 0.134 / 0.077	+2.8 percentage points; weakest low-VAF discrimination and calibration among the evaluated models.




Fig 7 Radar Comparison of Multimodal Biomarker Contribution, Calibration Reliability, and Low-Abundance ctDNA Detection

Figure 7 shows that LMCS-Net maintains the broadest performance envelope across all six normalized surveillance dimensions. Low-VAF sensitivity reached 88.7%, compared with 81.6% for the transformer, 79.8% for LSTM, 78.9% for GRU, 75.4% for XGBoost, and 72.6% for random forest. Low-VAF specificity was likewise highest at 92.1%, while AUPRC reached 93.1% on the percentage scale. Calibration was strongest for LMCS-Net, with Brier reliability of 93.2% and ECE reliability of 97.6%, derived from raw Brier and ECE values of 0.068 and 0.024. The multimodal contribution index was normalized to the largest observed AUPRC gain: LMCS-Net therefore scored 100, versus 59.3 for the transformer, 50.0 for LSTM, 46.5 for GRU, 37.2 for XGBoost, and 32.6 for random forest. The pattern indicates that multimodal molecular fusion improves sensitivity without sacrificing probability reliability at low ctDNA abundance. This separation is particularly important when molecular progression signals are weak, heterogeneous, and distributed across several biomarker classes simultaneously.

For Figure 7, Brier reliability = $(100(1 - \text{Brier score}))$, ECE reliability = $(100(1 - \text{ECE}))$, and the multimodal contribution index expresses each algorithm's

AUPRC gain relative to the maximum observed gain of 8.6 percentage points.

➤ Ablation, Sensitivity, Robustness, and Patient-Level Interpretability Analysis

Ablation and perturbation analyses examined whether LMCS-Net performance depended materially on its proposed architectural components and whether predictions remained stable under degraded molecular information. Table 5 shows that the complete LMCS-Net achieved an AUROC of 0.947, compared with 0.901 when restricted to ctDNA-only information, confirming a measurable contribution from multimodal biomarker integration. Under 20% random biomarker missingness and 10% measurement noise, AUROC remained 0.925 and 0.918, respectively, indicating comparatively strong robustness. Patient-level attribution stability reached 0.910, exceeding all comparator algorithms. Targeted LMCS-Net ablation further reduced AUROC to 0.918 without temporal attention, 0.913 without gated multimodal fusion, 0.921 without patient-specific normalization, and 0.928 without fragmentomic inputs. Removing Bayesian calibration increased the Brier score from 0.068 to 0.093, demonstrating that uncertainty calibration primarily improved probability reliability rather than ranking discrimination.

Table 5 Comparative Ablation, Robustness, Sensitivity, and Patient-Level Interpretability Performance

Algorithm	Full / ctDNA-Only AUROC	Robustness and Interpretability Metrics	Interpretation
LMCS-Net	0.947 / 0.901	Missing-data AUROC 0.925; noise AUROC 0.918; attribution stability 0.910	Largest multimodal benefit while retaining the strongest robustness and most stable patient-level feature attribution.
Transformer	0.914 / 0.880	Missing-data AUROC 0.881; noise AUROC 0.873; attribution stability 0.840	Strong temporal modeling but greater deterioration when molecular modalities are removed or corrupted.
LSTM	0.901 / 0.872	Missing-data AUROC 0.867; noise AUROC 0.858; attribution stability 0.810	Maintains useful sequential information but shows reduced robustness to incomplete longitudinal biomarker profiles.
GRU	0.894 / 0.866	Missing-data AUROC 0.859; noise AUROC 0.850; attribution stability 0.790	Computationally efficient recurrent baseline with moderate sensitivity to perturbation and weaker explanation stability.
XGBoost	0.879 / 0.854	Missing-data AUROC 0.836; noise AUROC 0.828; attribution stability 0.730	Nonlinear feature interactions remain useful, but performance declines substantially under missing or noisy biomarkers.
Random Forest	0.861 / 0.838	Missing-data AUROC 0.817; noise AUROC 0.809; attribution stability 0.690	Lowest robustness and patient-level attribution consistency among the evaluated models.




Fig 8 Scatter-Profile Comparison of Ablation, Missing-Biomarker Robustness, Measurement-Noise Sensitivity, and Patient-Level Interpretability

Figure 8 shows persistent separation between LMCS-Net and the five benchmark algorithms across all robustness conditions. Full-model AUROC was 94.7% for LMCS-Net, decreasing to 90.1% under ctDNA-only ablation, a 4.6 percentage-point loss that quantifies the contribution of complementary molecular biomarkers. With 20% biomarker missingness, LMCS-Net retained 92.5%, compared with 88.1% for the transformer, 86.7% for LSTM, 85.9% for GRU, 83.6% for XGBoost, and 81.7% for random forest. Under 10% measurement noise, corresponding performance remained 91.8%, 87.3%, 85.8%, 85.0%, 82.8%, and 80.9%, respectively. Patient-level attribution stability was also highest for LMCS-Net at 91.0%, compared with 84.0% for the transformer and 69.0% for random forest. The combined pattern indicates that temporal attention, individualized normalization, fragmentomic information, multimodal fusion, and uncertainty-aware modeling contribute not only to predictive accuracy but also to resilience and interpretable patient-specific surveillance.

V. CONCLUSIONS AND RECOMMENDATIONS

➤ Summary of Major Findings and Contributions of LMCS-Net

The study developed LMCS-Net as an integrated longitudinal cancer-surveillance architecture that combines circulating tumor DNA dynamics with fragmentomic, methylation, copy-number, somatic-mutation, protein-biomarker, treatment-exposure, and sampling-interval information. The principal finding was that jointly modeling temporal molecular trajectories provided stronger surveillance performance than interpreting isolated ctDNA measurements or applying static machine-learning models. LMCS-Net achieved an AUROC of 0.947, AUPRC of 0.931, sensitivity of 0.912,

specificity of 0.924, and F1-score of 0.918, consistently outperforming transformer, LSTM, GRU, XGBoost, and random-forest comparators. Longitudinal progression sensitivity reached 91.2%, recurrence sensitivity reached 89.6%, and the model produced a median molecular lead time of 4.8 months while maintaining a false-alert rate of 5.8%. These results indicate that temporally persistent molecular changes can be identified substantially earlier than conventional confirmation without disproportionately increasing false-positive surveillance signals.

A second major contribution was the explicit treatment of low-abundance and heterogeneous molecular information. LMCS-Net achieved 88.7% sensitivity and 92.1% specificity for low-VAF detection, supported by adaptive reliability weighting, individualized baseline normalization, temporal attention, and gated multimodal fusion. Its Brier score of 0.068 and expected calibration error of 0.024 demonstrated strong probability calibration. Ablation analysis further showed that removing temporal attention, multimodal fusion, normalization, or fragmentomic information reduced performance, while eliminating Bayesian calibration primarily degraded probability reliability. Collectively, the findings establish LMCS-Net as a unified computational framework for transforming serial molecular measurements into patient-specific, uncertainty-aware surveillance trajectories suitable for monitoring treatment response, residual disease, molecular progression, and recurrence.

➤ Clinical and Translational Implications for Longitudinal Precision Oncology Monitoring

The clinical relevance of LMCS-Net lies in its capacity to shift molecular surveillance from episodic threshold-based interpretation toward continuously updated patient-specific risk assessment. In conventional monitoring, a patient may be classified as ctDNA positive

or negative according to a fixed analytical cutoff even when the clinically important information is contained in the trajectory, persistence, or concordance of several weak molecular signals. LMCS-Net instead evaluates how ctDNA abundance, mutant molecules, tumor fraction, fragmentomic abnormalities, methylation patterns, copy-number alterations, mutations, protein biomarkers, and treatment exposure evolve jointly over time. This permits a gradual but persistent molecular rise to be recognized as potentially significant even when individual measurements remain near conventional detection thresholds. The 4.8-month median molecular lead time observed in the study illustrates the potential translational value of this approach for initiating confirmatory imaging, intensified surveillance, or multidisciplinary review before overt clinical progression becomes apparent.

The model also supports more individualized interpretation of treatment response. A falling risk trajectory accompanied by decreasing ctDNA fraction and stabilizing complementary biomarkers could indicate molecular response, whereas a sustained increase in risk associated with emerging mutations and concordant fragmentomic or copy-number changes could indicate treatment resistance. Importantly, Bayesian uncertainty estimates and calibrated probabilities provide additional context for such signals. A high-risk prediction accompanied by substantial uncertainty should not be interpreted identically to a similarly high probability supported by stable, concordant molecular evidence. This distinction is essential for clinical translation because surveillance tools should support rather than replace oncological decision-making. Integration into precision-oncology workflows would therefore require clearly defined alert thresholds, explainable feature contributions, laboratory-quality controls, interoperable reporting, and prospective evaluation of whether earlier molecular alerts improve clinically meaningful outcomes.

➤ *Study Limitations, Model Generalizability, and Validation Requirements*

Several limitations must be addressed before LMCS-Net can be regarded as suitable for routine clinical deployment. First, longitudinal molecular data are intrinsically heterogeneous. ctDNA shedding differs according to cancer type, disease burden, metastatic site, vascularity, treatment mechanism, and biological turnover. A surveillance model developed from one tumor population may therefore perform differently in cancers characterized by low plasma shedding or distinct genomic architectures. Second, molecular assays vary in sequencing depth, panel design, limit of detection, sample-processing workflow, methylation coverage, fragmentomic extraction, and copy-number resolution. These platform-specific differences may shift feature distributions and impair transportability if model normalization and calibration are not reassessed across institutions. Irregular sampling intervals, missing biomarker modalities, and incomplete matched leukocyte sequencing could further affect interpretation of low-abundance variants and clonal-hematopoiesis confounding.

The robustness analyses indicate that LMCS-Net retains comparatively strong performance under 20% missing biomarkers and 10% simulated measurement noise, but these perturbation tests cannot substitute for external clinical validation. Generalizability should therefore be evaluated using geographically, demographically, technologically, and clinically independent cohorts. Validation should include multiple cancer types, treatment regimens, disease stages, sequencing platforms, and laboratory environments. Calibration must be reassessed prospectively because probability reliability can deteriorate even when discrimination remains acceptable. Clinical validation should also determine whether the 4.8-month lead-time advantage remains stable across tumor types and whether earlier molecular alerts produce actionable information rather than unnecessary investigations. Future evaluation should incorporate prospective locked-model testing, predefined endpoints, decision-curve analysis, time-dependent discrimination, calibration-in-the-large, subgroup analyses, and comparison with standard imaging and biomarker pathways. Regulatory-grade deployment would additionally require reproducibility, auditability, cybersecurity, version control, and continuous post-deployment performance monitoring.

➤ *Recommendations for Future Multimodal and Data-Driven Cancer Surveillance Systems*

Future cancer-surveillance systems should extend beyond single-analyte ctDNA monitoring toward multimodal longitudinal architectures that explicitly represent biological heterogeneity, measurement uncertainty, treatment context, and patient-specific molecular baselines. Priority should be given to prospective multicenter studies in which ctDNA variants, mutant molecules per milliliter, tumor fraction, fragmentomic signatures, methylation profiles, copy-number alterations, protein biomarkers, imaging findings, treatment exposure, and clinical outcomes are collected according to harmonized schedules. Such datasets would permit rigorous assessment of whether the multimodal gains demonstrated by LMCS-Net persist across disease sites, therapies, populations, and analytical platforms. Future models should also incorporate missing-data-aware learning so that clinically useful predictions remain possible when particular assays are unavailable, delayed, or technically unsuccessful.

Methodologically, development should focus on dynamic calibration, causal treatment-response modeling, uncertainty decomposition, and interpretable temporal attribution. Distinguishing epistemic uncertainty caused by insufficient training data from aleatoric uncertainty caused by biological variability would improve confidence-aware clinical reporting. Patient-specific digital trajectories could additionally be used to simulate alternative surveillance intervals and determine when additional blood sampling is likely to provide meaningful information. Federated or privacy-preserving learning should be considered for multicenter model development where direct exchange of genomic data is restricted. Prospective trials should evaluate whether model-guided

surveillance improves recurrence detection, treatment adaptation, quality of life, healthcare utilization, and survival rather than reporting predictive accuracy alone. Thresholds for clinical alerts should be determined according to disease context and consequence of false-positive or false-negative decisions. Ultimately, future systems should combine strong discrimination with calibration, transparent biomarker attribution, standardized laboratory inputs, secure data governance, external validation, and clinically interpretable outputs so that longitudinal molecular surveillance can function as a reliable component of precision-oncology decision support.

REFERENCES

- [1]. Abbosh, C., Birkbak, N. J., Wilson, G. A., Jamal-Hanjani, M., Constantin, T., Salari, R., Le Quesne, J., Moore, D. A., Veeriah, S., Rosenthal, R., Marafioti, T., Kirkizlar, E., Watkins, T. B. K., McGranahan, N., Ward, S., Martinson, L., Riley, J., Fraioli, F., Al Bakir, M., . . . Swanton, C. (2017). Phylogenetic ctDNA analysis depicts early-stage lung cancer evolution. *Nature*, 545(7655), 446–451. DOI: 10.1038/nature22364.
- [2]. Assaf, Z. J. F., Zou, W., Fine, A. D., Socinski, M. A., Young, A., Lipson, D., Freidin, J. F., Kennedy, M., Polisecki, E., Nishio, M., Fabrizio, D., Oxnard, G. R., Cummings, C., Rode, A., Reck, M., Patil, N. S., Lee, M., Shames, D. S., & Schulze, K. (2023). A longitudinal circulating tumor DNA-based model associated with survival in metastatic non-small-cell lung cancer. *Nature Medicine*, 29(4), 859–868. <https://doi.org/10.1038/s41591-023-02226-6>
- [3]. Atalor, S. I. (2019). Federated learning architectures for predicting adverse drug events in oncology without compromising patient privacy. *Iconic Research and Engineering Journals*, 2(12), 246–269.
- [4]. Atalor, S. I. (2022). Blockchain-enabled pharmacovigilance infrastructure for national cancer registries. *International Journal of Scientific Research and Modern Technology*, 1(1), 50–64. DOI: 10.38124/ijrsmt.v1i1.493.
- [5]. Atalor, S. I. (2022). Data-driven cheminformatics models for predicting bioactivity of natural compounds in oncology. *International Journal of Scientific Research and Modern Technology*, 1(1), 65–76. DOI: 10.38124/ijrsmt.v1i1.496.
- [6]. Avoseh, N. R., & Enyejo, J. O. (2022). Agentic artificial intelligence for enterprise data governance and decision support through secure retrieval augmented generation across Microsoft 365 information environments. *International Journal of Multidisciplinary Futuristic Development*, 3(1), 65–82. DOI: 10.54660/IJMFDF.2022.3.1.65-82.
- [7]. Boehm, K. M., Khosravi, P., Vanguri, R., Gao, J., & Shah, S. P. (2022). Harnessing multimodal data integration to advance precision oncology. *Nature Reviews Cancer*, 22(2), 114–126. <https://doi.org/10.1038/s41568-021-00408-3>
- [8]. Bruhm, D. C., Vulpescu, N. A., Foda, Z. H., Phallen, J., Scharpf, R. B., & Velculescu, V. E. (2025). Genomic and fragmentomic landscapes of cell-free DNA for early cancer detection. *Nature Reviews Cancer*, 25(5), 341–358. <https://doi.org/10.1038/s41568-025-00795-x>
- [9]. Chabon, J. J., Hamilton, E. G., Kurtz, D. M., Esfahani, M. S., Moding, E. J., Stehr, H., Schroers-Martin, J., Nabet, B. Y., Chen, B., Chaudhuri, A. A., Liu, C. L., Hui, A. B., Jin, M. C., Azad, T. D., Almanza, D., Jeon, Y.-J., Nesselbush, M. C., Keh, L. C. T., Bonilla, R. F., . . . Diehn, M. (2020). Integrating genomic features for non-invasive early lung cancer detection. *Nature*, 580(7802), 245–251. DOI: 10.1038/s41586-020-2140-0.
- [10]. Cheerla, A., & Gevaert, O. (2019). Deep learning with multimodal representation for pancancer prognosis prediction. *Bioinformatics*, 35(14), i446–i454. DOI: 10.1093/bioinformatics/btz342.
- [11]. Cohen, J. D., Li, L., Wang, Y., Thoburn, C., Afsari, B., Danilova, L., Douville, C., Javed, A. A., Wong, F., Mattox, A., Hruban, R. H., Wolfgang, C. L., Goggins, M. G., Dal Molin, M., Wang, T.-L., Roden, R., Klein, A. P., Ptak, J., Dobbyn, L., . . . Papadopoulos, N. (2018). Detection and localization of surgically resectable cancers with a multi-analyte blood test. *Science*, 359(6378), 926–930. DOI: 10.1126/science.aar3247.
- [12]. Corcoran, R. B., & Chabner, B. A. (2018). Application of cell-free DNA analysis to cancer treatment. *The New England Journal of Medicine*, 379(18), 1754–1765. <https://doi.org/10.1056/NEJMr1706174>
- [13]. Cristiano, S., Leal, A., Phallen, J., Fiksel, J., Adleff, V., Bruhm, D. C., Jensen, S. Ø., Medina, J. E., Hruban, C., White, J. R., Palsgrove, D. N., Niknafs, N., Anagnostou, V., Forde, P., Naidoo, J., Marrone, K., Brahmer, J., Woodward, B. D., Husain, H., . . . Velculescu, V. E. (2019). Genome-wide cell-free DNA fragmentation in patients with cancer. *Nature*, 570(7761), 385–389. <https://doi.org/10.1038/s41586-019-1272-6>
- [14]. Dawson, S.-J., Tsui, D. W. Y., Murtaza, M., Biggs, H., Rueda, O. M., Chin, S.-F., Dunning, M. J., Gale, D., Forshew, T., Mahler-Araujo, B., Rajan, S., Humphray, S., Becq, J., Halsall, D., Wallis, M., Bentley, D., Caldas, C., & Rosenfeld, N. (2013). Analysis of circulating tumor DNA to monitor metastatic breast cancer. *The New England Journal of Medicine*, 368(13), 1199–1209. <https://doi.org/10.1056/NEJMoa1213261>
- [15]. Diaz, L. A., Jr., & Bardelli, A. (2014). Liquid biopsies: Genotyping circulating tumor DNA. *Journal of Clinical Oncology*, 32(6), 579–586. DOI: 10.1200/JCO.2012.45.2011.
- [16]. Diehl, F., Schmidt, K., Choti, M. A., Romans, K., Goodman, S., Li, M., Thornton, K., Agrawal, N., Sokoll, L., Szabo, S. A., Kinzler, K. W., Vogelstein, B., & Diaz, L. A., Jr. (2008). Circulating mutant DNA to assess tumor dynamics. *Nature Medicine*, 14(9), 985–990. <https://doi.org/10.1038/nm.1789>

- [17]. Heitzer, E., Haque, I. S., Roberts, C. E. S., & Speicher, M. R. (2019). Current and future perspectives of liquid biopsies in genomics-driven oncology. *Nature Reviews Genetics*, 20, 71–88. DOI: 10.1038/s41576-018-0071-5.
- [18]. Idika, C. N. (2022). Lightweight authentication mechanisms for securing wearable medical devices in body area networks. *Global Journal of Multidisciplinary Studies*, 11(9), 75–89. DOI: 10.5281/zenodo.17519630.
- [19]. Idika, C. N., Salami, E. O., Ijiga, O. M., & Enyejo, L. A. (2021). Deep learning driven malware classification for cloud-native microservices in edge computing architectures. *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, 7(4), 710–732. DOI: 10.32628/CSEIT182551.
- [20]. Ignatiadis, M., Sledge, G. W., & Jeffrey, S. S. (2021). Liquid biopsy enters the clinic: Implementation issues and future challenges. *Nature Reviews Clinical Oncology*, 18(5), 297–312. DOI: 10.1038/s41571-020-00457-x.
- [21]. Ijiga, O. M., Ifenatuora, G. P., & Olateju, M. (2022). AI-powered e-learning platforms for STEM education: Evaluating effectiveness in low-bandwidth and remote learning environments. *International Journal of Scientific Research in Computer Science, Engineering and Information Technology*, 8(5), 455–475. DOI: 10.32628/CSEIT23902187.
- [22]. Imoh, P. O., & Idoko, I. P. (2022). Gene-environment interactions and epigenetic regulation in autism etiology through multi-omics integration and computational biology approaches. *International Journal of Scientific Research and Modern Technology*, 1(8), 1–16. DOI: 10.38124/ijrsmt.v1i8.463.
- [23]. Kurtz, D. M., Esfahani, M. S., Scherer, F., Soo, J., Jin, M. C., Liu, C. L., Newman, A. M., Dührsen, U., Hüttmann, A., Casasnovas, O., Westin, J. R., Ritgen, M., Böttcher, S., Langerak, A. W., Roschewski, M., Wilson, W. H., Gaidano, G., Rossi, D., Bahlo, J., . . . Alizadeh, A. A. (2019). Dynamic risk profiling using serial tumor biomarkers for personalized outcome prediction. *Cell*, 178(3), 699–713.e19. <https://doi.org/10.1016/j.cell.2019.06.011>
- [24]. Kwarteng, R. A., Idoko, I. P., & Ijiga, O. M. (2021). Data-driven project management frameworks for improving IT service delivery in distributed organizations. *Computer Science & IT Research Journal*, 2(1), 76–105. DOI: 10.51594/csitrj.v2i1.2172.
- [25]. Liu, M. C., Oxnard, G. R., Klein, E. A., Swanton, C., Seiden, M. V., & CCGA Consortium. (2020). Sensitive and specific multi-cancer detection and localization using methylation signatures in cell-free DNA. *Annals of Oncology*, 31(6), 745–759. <https://doi.org/10.1016/j.annonc.2020.02.011>
- [26]. Merker, J. D., Oxnard, G. R., Compton, C., Diehn, M., Hurley, P., Lazar, A. J., Lindeman, N., Lockwood, C. M., Rai, A. J., Schilsky, R. L., Tsimberidou, A. M., Vasalos, P., Billman, B. L., Oliver, T. K., Bruinooge, S. S., Hayes, D. F., & Turner, N. C. (2018). Circulating tumor DNA analysis in patients with cancer: American Society of Clinical Oncology and College of American Pathologists joint review. *Journal of Clinical Oncology*, 36(16), 1631–1641. DOI: 10.1200/JCO.2017.76.8671.
- [27]. Mobadersany, P., Yousefi, S., Amgad, M., Gutman, D. A., Barnholtz-Sloan, J. S., Velázquez Vega, J. E., Brat, D. J., & Cooper, L. A. D. (2018). Predicting cancer outcomes from histology and genomics using convolutional networks. *Proceedings of the National Academy of Sciences of the United States of America*, 115(13), E2970–E2979. DOI: 10.1073/pnas.1717139115.
- [28]. Newman, A. M., Lovejoy, A. F., Klass, D. M., Kurtz, D. M., Chabon, J. J., Scherer, F., Stehr, H., Liu, C. L., Bratman, S. V., Say, C., Zhou, L., Carter, J. N., West, R. B., Sledge, G. W., Jr., Shrager, J. B., Loo, B. W., Jr., Neal, J. W., Wakelee, H. A., Diehn, M., & Alizadeh, A. A. (2016). Integrated digital error suppression for improved detection of circulating tumor DNA. *Nature Biotechnology*, 34(5), 547–555. <https://doi.org/10.1038/nbt.3520>
- [29]. Nwokocha, C. R., Soetan, K. T., & Peter-Anyebe, A. C. (2022). Automating ETL pipelines with SQL Server Integration Services to improve health data quality and reporting accuracy in national health systems. *International Journal of Scientific Research in Science and Technology*, 9(4), 804–827.
- [30]. Ocharo, D. O., & Omachi, A. (2022). Optimization of microgrid-controlled chiller plants for data center cooling in the Northeastern United States. *International Journal of Scientific Research in Science and Technology*, 9(3), 865–880.
- [31]. Onyekaonwu, C. B., Peter-Anyebe, A. C., & Raphael, F. O. (2019). From prescription to prediction: Leveraging AI/ML to improve medication adherence and adverse drug event detection in community pharmacies. *International Journal of Scientific Research in Science and Technology*, 6(5), 460–476.
- [32]. Razavi, P., Li, B. T., Brown, D. N., Jung, B., Hubbell, E., Shen, R., Abida, W., Juluru, K., De Bruijn, I., Hou, C., Venn, O., Lim, R., Anand, A., Maddala, T., Gnerre, S., Satya, R. V., Liu, Q., Shen, L., Eattock, N., . . . Reis-Filho, J. S. (2019). High-intensity sequencing reveals the sources of plasma circulating cell-free DNA variants. *Nature Medicine*, 25(12), 1928–1937. <https://doi.org/10.1038/s41591-019-0652-7>
- [33]. Reinert, T., Henriksen, T. V., Christensen, E., Sharma, S., Salari, R., Sethi, H., Knudsen, M., Nordentoft, I., Wu, H.-T., Tin, A. S., Rasmussen, M. H., Vang, S., Shchegrova, S., Johansen, A. F. B., Srinivasan, R., Assaf, Z., Balcioglu, M., Olson, A., Dashner, S., . . . Andersen, C. L. (2019). Analysis of plasma cell-free DNA by ultradeep sequencing in patients with stages I to III colorectal cancer. *JAMA Oncology*, 5(8), 1124–1131. DOI: 10.1001/jamaoncol.2019.0528.

- [34]. Shen, S. Y., Singhania, R., Fehringer, G., Chakravarthy, A., Rochrl, M. H. A., Chadwick, D., Zuzarte, P. C., Borgida, A., Wang, T. T., Li, T., Kis, O., Zhao, Z., Spreafico, A., da Silva Medina, T., Wang, Y., Roulois, D., Ettayebi, I., Chen, Z., Chow, S., . . . De Carvalho, D. D. (2018). Sensitive tumour detection and classification using plasma cell-free DNA methylomes. *Nature*, *563*(7732), 579–583. DOI: 10.1038/s41586-018-0703-0.
- [35]. Snyder, M. W., Kircher, M., Hill, A. J., Daza, R. M., & Shendure, J. (2016). Cell-free DNA comprises an in vivo nucleosome footprint that informs its tissues-of-origin. *Cell*, *164*(1–2), 57–68. DOI: 10.1016/j.cell.2015.11.050.
- [36]. Wan, J. C. M., Massie, C., Garcia-Corbacho, J., Mouliere, F., Brenton, J. D., Caldas, C., Pacey, S., Baird, R., & Rosenfeld, N. (2017). Liquid biopsies come of age: Towards implementation of circulating tumour DNA. *Nature Reviews Cancer*, *17*(4), 223–238. <https://doi.org/10.1038/nrc.2017.7>