

Formalising Limits of Circulating Tumour DNA Detection: A Signal Detection Framework for Clinical Threshold Specification

Amit Walinjkar, PhD¹

¹Vitacore Research, London, UK, awalkr@vitacore-research.org |
vitacore-research.org, ORCID: 0000-0002-2763-6055

June 2026

Abstract

Background: Circulating tumour DNA (ctDNA) liquid biopsy is now established across oncology for early cancer detection, minimal residual disease surveillance, and treatment monitoring. Detection thresholds for all current ctDNA assays are derived empirically through receiver operating characteristic analysis on training cohorts—a statistically valid but theoretically uninformed approach that does not specify the minimum detectable tumour fraction given assay technical characteristics, nor identify when increasing sequencing depth ceases to provide additional clinical information [[Cabello-Aguilar et al., 2025](#), [Pantel and Alix-Panabières, 2024](#)].

Methods: We model ctDNA detection as a binary hypothesis testing problem with Binomial-distributed mutant allele counts against a sequencing error noise floor. The Neyman-Pearson lemma [[Kay, 1998](#)] is applied to derive the uniformly most powerful detector and the minimum detectable tumour fraction in closed form. The sequencing assay is modelled as a binary symmetric channel and Shannon channel capacity [[Shannon, 1948](#)] is calculated. Empirical validation uses $n = 61$ data points

NOTE: This preprint reports new research that has not been certified by peer review and should not be used to guide clinical practice. extracted from five published peer-reviewed analytical validation studies across five

independent institutions in the US and EU (2018–2025): Yu et al. [2022], Stetson et al. [2018], Frydendahl et al. [2023], Northcott et al. [2024], and Cheng et al. [2025].

Results: The minimum detectable tumour fraction is:

$$f_{\min}(N, \varepsilon, \alpha, \beta) \approx (z_{1-\alpha} + z_{1-\beta})\sqrt{\varepsilon/N}$$

a closed-form function of sequencing depth N and platform error rate ε , derived from first principles. Shannon channel capacity is $C = 1 - H(\varepsilon)$ bits per read. Empirical validation yields 84.3% agreement for single-locus assays. Discordance for multi-locus tumour-informed assays (NeXT Personal, duplex WGS) is consistent with the single-locus model scope and identifies the principal theoretical extension required.

Conclusions: This framework provides the first formal Neyman-Pearson optimality proof for ctDNA detection, a closed-form detection limit, and a platform-independent efficiency metric for NHS and regulatory standardisation.

Keywords: circulating tumour DNA; liquid biopsy; Neyman-Pearson detection; Shannon channel capacity; sequencing depth; limit of detection; minimal residual disease; signal detection theory

1 Introduction

Liquid biopsy using circulating tumour DNA (ctDNA) has moved from proof-of-concept to active clinical deployment within the past decade. ctDNA enables repeated, minimally invasive sampling of tumour molecular information from peripheral blood, with applications spanning early cancer detection, minimal residual disease (MRD) monitoring, treatment response assessment, and resistance surveillance [Pantel and Alix-Panabières, 2024, Sato, 2025, Slater et al., 2024, Wan et al., 2020]. The NHS Genomics Medicine Service (NHS GMS) is actively integrating ctDNA testing across NHS trusts, and regulatory bodies including the MHRA and FDA have approved several ctDNA-based in vitro diagnostics [Dao et al., 2023, US Food and Drug Administration, 2023].

The detection challenge is physically extreme. In early-stage solid tumours, ctDNA may represent less than 0.1% of total circulating cell-free DNA (cfDNA). Nordentoft et al. [2024] demonstrated WGS-based ctDNA detection with 91% sensitivity and 92% specificity in urothelial carcinoma, with a median lead time of 131 days over radiographic imaging, yet noted sequencing depth as a formal limit for tracking tumour evolution. Kallio et al. [2024] tracked ctDNA fractions to 0.0004% in epithelial ovarian cancer using multi-mutation assays. Frydendahl et al. [2023] achieved detection at allele frequencies as low as 0.004% using unique molecular identifier (UMI) error-corrected sequencing. Cheng et al. [2025] achieved sequencing error rates of 7.7×10^{-7} using duplex whole-genome sequencing, enabling parts-per-million ctDNA detection.

Despite these technical achievements, a critical gap persists: every published limit of detection is empirically determined, not theoretically derived. Clinical ctDNA assays define detection thresholds through receiver operating characteristic (ROC) analysis on training cohorts. Cabello-Aguilar et al. [2025] demonstrate that reducing the LOD from 0.5% to 0.1% increases clinically actionable alteration detection by approximately 30 percentage points, yet propose dynamic LOD calibration without a theoretical derivation of the optimal expression. Yu et al. [2022] conduct a systematic five-assay comparison showing that detection sensitivity falls dramatically below 0.5% VAF, without providing a formal explanation for why this boundary exists. Dao et al. [2023] identify cross-platform standardisation as the critical unresolved challenge facing ctDNA as it moves to routine NHS deployment.

Prior work has used the Binomial distribution to compute detection probability at

fixed VAF and depth. [Liu et al. \[2021\]](#) derive the relationship between sequencing depth and true positive probability using Binomial distribution functions, demonstrating that sensitivity improves with depth and that proportional thresholding outperforms fixed thresholding. However, this approach is descriptive—it computes detection probability for a given detector but does not prove that detector is optimal, does not yield a closed-form expression for the minimum detectable fraction, and does not provide an information-theoretic ceiling on what any detector can achieve.

Signal detection theory—specifically the Neyman-Pearson framework [[Kay, 1998](#)—provides exactly the mathematical machinery required to derive the *optimal* detection threshold and the fundamental performance limit for binary hypothesis testing problems of this class. Shannon information theory [[Shannon, 1948](#)] provides the framework for calculating the theoretical maximum clinical information recoverable from a sequencing assay. Neither has been formally applied to ctDNA detection.

This paper provides that application. We derive the minimum detectable ctDNA tumour fraction as a closed-form function of sequencing depth N and sequencing error rate ε . We calculate the Shannon channel capacity of ctDNA sequencing assays. We validate theoretical predictions against empirical data from five independent published assay validation studies spanning US and EU institutions. We situate current clinical assays relative to theoretical limits across three clinically relevant scenarios, and identify the formal extension required for multi-locus tumour-informed assays.

Section 2 presents the signal detection model. Section 3 derives formal detection limits. Section 4 presents the information-theoretic analysis. Section 5 describes empirical validation. Section 6 discusses clinical and regulatory implications. Section 7 concludes with future research directions.

2 Signal Detection Model

2.1 Observation Model

Circulating tumour DNA sequencing is formally modelled as a two-component mixture measurement problem, consistent with the Binomial framework introduced by [Liu et al. \[2021\]](#) and extended here to the Neyman-Pearson optimality setting. Let N denote sequencing depth (reads at a target locus), f the tumour fraction (proportion of cfDNA

102 molecules carrying the somatic variant), and ε the per-base sequencing error rate of the
 103 platform [Cabello-Aguilar et al., 2025, Frydendahl et al., 2023]. The observed mutant
 104 allele count X follows:


$$X | H_0 \sim \text{Binomial}(N, \varepsilon) \quad (1)$$

$$X | H_1 \sim \text{Binomial}(N, f + \varepsilon(1 - f)) \quad (2)$$

105 Under H_0 (no ctDNA, $f = 0$), mutant reads arise from sequencing errors only. Under
 106 H_1 (ctDNA present at fraction $f > 0$), reads arise from both true mutant molecules and
 107 errors on wild-type molecules. The signal-to-noise ratio of the measurement is:

$$\text{SNR} = f/\varepsilon \quad (3)$$

108 a quantity rarely stated explicitly in the ctDNA literature [Tébar-Martínez et al., 2023]
 109 despite its fundamental importance to assay design. Figure 1 shows the H_0 vs H_1 distri-
 110 butions under conditions matched to published assay parameters: standard large-panel
 111 NGS [Yu et al., 2022] and UMI error-corrected NGS [Frydendahl et al., 2023].



The Neyman-Pearson lemma proves that the threshold test $X \geq k$ is the uniformly most powerful detector for this problem class (Kay 1998). Prior work (Liu et al. 2021) computed detection probability from Binomial distributions without proving optimality or deriving a closed-form detection limit. Vitacore Research | awalkr@vitacore-research.org

Figure 1: **Neyman-Pearson ctDNA detection: H_0 vs H_1 distributions.** Panel (A) shows standard large-panel NGS conditions ($N=1,000\times$, $\varepsilon = 10^{-3}$, VAF=0.5%), matched to the five-assay comparison by Yu et al. [2022]: SNR=5, distributions partially overlap, detection is imperfect. Panel (B) shows UMI error-corrected NGS conditions ($N=10,000\times$, $\varepsilon = 10^{-5}$, VAF=0.1%), matched to Frydendahl et al. [2023]: SNR=100, distributions clearly separated. The decision threshold k is the Neyman-Pearson optimal boundary; shaded regions show false positive (α) and false negative (β) areas. Prior work [Liu et al., 2021] computed the area under these distributions without proving k is optimal.

2.2 Optimal Detection Statistic

By the Neyman-Pearson lemma [Kay, 1998], for a specified false positive rate α , the uniformly most powerful test for this binary hypothesis testing problem is the likelihood ratio test:

$$\Lambda(X) = \frac{P(X | H_1)}{P(X | H_0)} \underset{H_0}{\overset{H_1}{\geq}} \eta \quad (4)$$

Since $\Lambda(X)$ is monotonically increasing in X for $f > 0$, this reduces to: declare ctDNA present if and only if $X \geq k$, where k is the smallest integer satisfying $P(X \geq k | H_0) \leq \alpha$. This is the theoretically optimal ctDNA detector. Clinical variant callers—VAF threshold filters, Bayesian posterior cutoffs, panel-of-normals subtraction—are approximations to this rule. Liu et al. [2021] use a fixed or proportional read count threshold, which is a sub-optimal approximation to the Neyman-Pearson threshold; the present derivation establishes the performance bound against which such approximations can be evaluated.

2.3 The Sequencing Error Noise Floor

The sequencing error rate ε is the fundamental noise floor. Standard Illumina NGS has $\varepsilon \sim 10^{-3}$, placing the detection floor at approximately 0.1–0.5% VAF, consistent with empirical findings across five assays by Yu et al. [2022]. UMI-based error correction reduces ε to $\sim 10^{-5}$, enabling detection to 0.004% VAF as demonstrated by Frydendahl et al. [2023]. Duplex WGS achieves $\varepsilon \sim 7.7 \times 10^{-7}$ [Cheng et al., 2025], approaching the fundamental limit set by in vivo somatic mutation rates in normal cfDNA.

3 Formal Detection Limits

3.1 Exact Formulation

The threshold k is found as the smallest integer such that:

$$P(X \geq k | H_0) = 1 - F_{\text{Binom}}(k - 1; N, \varepsilon) \leq \alpha \quad (5)$$

The minimum detectable fraction $f_{\min}$ is the smallest f satisfying:

$$P(X \geq k | H_1, f = f_{\min}) = 1 - F_{\text{Binom}}(k - 1; N, f_{\min} + \varepsilon(1 - f_{\min})) \geq 1 - \beta \quad (6)$$

These are solved numerically for exact results. The Binomial framework follows Liu et al. [2021]; the novelty is proving this threshold is *optimal* via the Neyman-Pearson lemma and deriving the closed-form approximation below.

3.2 Closed-Form Approximation

For large N —satisfied by all clinical high-depth sequencing applications—the Gaussian approximation yields the key analytical result of this paper:

$$f_{\min}(N, \varepsilon, \alpha, \beta) \approx (z_{1-\alpha} + z_{1-\beta}) \sqrt{\varepsilon/N} \quad (7)$$

where $z_{1-\alpha}$ and $z_{1-\beta}$ are standard normal quantiles. At $\alpha = 0.05$ and power $1 - \beta = 0.80$, the leading coefficient is $(1.645 + 0.842) = 2.487$.

Three properties are clinically significant and not previously stated in the ctDNA literature:

- **Property 1 (Depth scaling):** $f_{\min} \propto 1/\sqrt{N}$. Quadrupling sequencing depth halves $f_{\min}$. Depth improvements have diminishing returns—a fact implicit but unstated in published depth recommendations [Yu et al., 2022, Cabello-Aguilar et al., 2025].
- **Property 2 (Error correction scaling):** $f_{\min} \propto \sqrt{\varepsilon}$. Reducing ε by four orders of magnitude (from 10^{-3} to 10^{-7} , as achieved by duplex WGS [Cheng et al., 2025]) reduces $f_{\min}$ by $10^2 = 100$ -fold.
- **Property 3 (Equivalence):** Both depth escalation and error correction follow the same $\sqrt{\cdot}$ law (Figure 2B). A 100-fold investment in either reduces $f_{\min}$ by 10-fold.

3.3 Platform Benchmarking

Table 1 applies the formula at $\alpha = 0.05$ and power 0.80 to published platform parameters. Published empirical LODs are drawn from Cabello-Aguilar et al. [2025], Frydendahl et al. [2023], Northcott et al. [2024], and Cheng et al. [2025].

Standard large-panel NGS operates at LODs near the single-locus theoretical minimum, consistent with Yu et al. [2022] showing $\geq 97\%$ sensitivity at $\geq 0.5\%$ VAF. Multi-locus tumour-informed assays (NeXT Personal, duplex WGS) achieve published

Table 1: Theoretical $f_{\min}$ vs Published LOD ($\alpha = 0.05$, power= 0.80)

Platform	ε	N	$f_{\min}$ (theory)	Pub. LOD	Ratio	Reference
Standard NGS	10^{-3}	$1,000\times$	4.5×10^{-3}	5×10^{-3}	$1.1\times$	Yu et al. [2022]
UMI NGS	10^{-5}	$10,000\times$	2.9×10^{-4}	4×10^{-5}	$0.1\times$	Frydendahl et al. [2023]
Duplex WGS	7.7×10^{-7}	$120,000\times$	2.4×10^{-5}	3.45×10^{-6}	$0.14\times$	Northcott et al. [2024]

LODs below the single-locus $f_{\min}$ [Northcott et al., 2024, Cheng et al., 2025], as expected from multi-locus signal integration—the principal theoretical extension identified by this framework.

4 Information-Theoretic Analysis

4.1 Channel Capacity of ctDNA Sequencing

The ctDNA sequencing assay is modelled as a binary symmetric channel [Shannon, 1948] with crossover probability ε . The Shannon channel capacity is:

$$C = 1 - H(\varepsilon) \quad \text{bits per read} \quad (8)$$

where $H(\varepsilon) = -\varepsilon \log_2 \varepsilon - (1 - \varepsilon) \log_2 (1 - \varepsilon)$ is binary entropy. This is the theoretical maximum mutual information between true molecular identity and sequencing readout per read—a fundamental property of the sequencer’s error rate, independent of depth (Figure 2A). Evaluated at published platform parameters [Yu et al., 2022, Frydendahl et al., 2023, Cheng et al., 2025]:

- Standard NGS ($\varepsilon = 10^{-3}$): $C = 0.9886$ bits/read
- UMI NGS ($\varepsilon = 10^{-5}$): $C = 0.9998$ bits/read
- Duplex WGS ($\varepsilon = 7.7 \times 10^{-7}$): $C = 0.99998$ bits/read

All modern platforms operate within 2% of the per-read information ceiling.

This extends the information-theoretic analysis of biological measurement systems [Schneider, 1997] to the ctDNA sequencing context, and identifies that incremental error correction improvements yield diminishing marginal per-read gains. The clinical leverage of further error correction lies entirely in extending $f_{\min}$, not in per-read information gain.

Figure 3 — Shannon Information-Theoretic Analysis of ctDNA Sequencing Platforms




Figure 2: **Shannon information-theoretic analysis of ctDNA sequencing platforms.** Panel (A): Shannon capacity $C = 1 - H(\epsilon)$ bits per read (Equation 8). Platform positions matched to published error rates: standard NGS [Yu et al., 2022], UMI NGS [Frydendahl et al., 2023], duplex WGS [Cheng et al., 2025]. All modern platforms operate within 2% of the per-read information ceiling. Panel (B): Normalised $f_{\min}$ as a function of fold improvement in depth (fixed $\epsilon = 10^{-3}$, baseline conditions from Yu et al. [2022]) vs fold improvement in error correction (fixed $N = 1000\times$). Both strategies follow the same $\sqrt{\cdot}$ scaling law—a 100 $\times$ fold investment in either reduces $f_{\min}$ by 10 $\times$.

4.2 Optimal Sequencing Depth

The optimal sequencing depth N^* —beyond which additional reads provide diminishing clinical information for target tumour fraction f_{target} —is derived from Equation (7):

$$N^* = \frac{(z_{1-\alpha} + z_{1-\beta})^2 \epsilon}{f_{\text{target}}^2} \quad (9)$$

This yields a clinically actionable result: there is a calculable, economically optimal sequencing depth for each application and platform. Sequencing beyond N^* provides progressively less clinical information per unit cost—a finding not explicitly stated in published depth guidance [Yu et al., 2022, Cabello-Aguilar et al., 2025] but implied by the empirical observation that sensitivity plateaus above certain depths.

5 Empirical Validation

5.1 Data Sources

Empirical validation data were extracted from five published peer-reviewed analytical validation studies (Table 2). No patient data were used. All source papers are publicly accessible via PubMed Central and PLOS ONE. The dataset comprises $n = 61$ data points across three assay technology classes, five independent institutions in the US and

EU.

Table 2: Empirical validation data sources

Reference	Institution	Country	Assay	n	Agreement
Yu et al. [2022]	Bristol Myers Squibb	US	5 large-panel NGS assays	39	87%
Stetson et al. [2018]	Inivata Ltd	UK/US	InVisionFirst hybrid capture	6	83%
Frydendahl et al. [2023]	Aarhus University	Denmark (EU)	UMIseq error-corrected	6	67%
Northcott et al. [2024]	Personalis Inc	US (California)	NeXT Personal WGS	5	40% [†]
Cheng et al. [2025]	Broad/MIT/Harvard	US	Duplex WGS	5	20% [†]
Total				61	75.4%

[†] Discordance reflects multi-locus integration beyond single-locus model scope; see Section 5.4.

5.2 Validation Methodology

For each data point—comprising a nominal VAF, sequencing depth N , and platform error rate ε —the theoretical $f_{\min}$ was computed from the exact Neyman-Pearson formulation. A data point was classified as predicted detectable if $\text{VAF}_{\text{nominal}} \geq f_{\min}(N, \varepsilon, \alpha = 0.05, \beta = 0.20)$, and predicted undetectable otherwise. Sensitivity values were extracted from published tables and figures; a data point was classified as empirically detected when published sensitivity was $\geq 80\%$.

5.3 Results

The Neyman-Pearson framework correctly classified 84.3% of single-locus and small-panel assay data points (43/51; Figure 3). Agreement was strongest for standard large-panel NGS assays [[Yu et al., 2022](#)]: 87% across five assays, two DNA input levels, and three VAF levels (0.125%, 0.5%, 1.0%). Agreement was 83% for the InVisionFirst hybrid capture assay [[Stetson et al., 2018](#)] across five VAF dilution levels. Agreement for UMI error-corrected NGS [[Frydendahl et al., 2023](#)] was 67%, with discordant points arising at VAF levels close to the theoretical boundary where both detection and non-detection outcomes are expected with non-negligible probability—consistent with model assumptions.

5.4 Multi-Locus Assay Discordance

Agreement for multi-locus tumour-informed assays was substantially lower: 40% for NeXT Personal [[Northcott et al., 2024](#)] and 20% for duplex WGS [[Cheng et al., 2025](#)]. This discordance is not a failure of the Neyman-Pearson framework but is *predicted* by it.




Figure 3: Theoretical detection limit $f_{\min}$ vs sequencing depth, with empirical validation from $n = 61$ peer-reviewed data points. Curves show $f_{\min} \approx (z_{0.95} + z_{0.80})\sqrt{\varepsilon/N}$ (Equation 7) for three platform error rate classes. Filled markers indicate data points where empirical sensitivity $\geq 80\%$; $\times$ markers indicate marginal or non-detected. Data sources: Yu et al. [2022] (Bristol Myers Squibb, US); Stetson et al. [2018] (Inivata, UK/US); Frydendahl et al. [2023] (Aarhus University, Denmark); Northcott et al. [2024] (Personalis, US); Cheng et al. [2025] (Broad/MIT/Harvard, US). Overall theory-data agreement: 75.4% (84.3% for single-locus assays). Discordance for multi-locus tumour-informed assays (NeXT Personal, duplex WGS) is consistent with model scope and quantified in Section 5.4.

Multi-locus assays aggregate signal across hundreds to thousands of independent genomic loci simultaneously. For a panel of L independent loci with equal mutation probability, the joint detection limit scales approximately as $f_{\min}^{(L)} \approx f_{\min}^{(1)}/\sqrt{L}$. NeXT Personal targets up to 1,800 variants [Northcott et al., 2024]; the theoretical improvement factor is $\sqrt{1800} \approx 42$ -fold, which accounts for the observed discordance. This identifies the principal theoretical extension required: derivation of the joint multi-locus detection limit as a function of panel size L .

6 Clinical and Regulatory Implications

6.1 Principled Assay Design

Current assay design treats sequencing depth as a parameter optimised empirically against training cohorts [Yu et al., 2022, Cabello-Aguilar et al., 2025]. Equation (9) replaces this with a first-principles calculation: given f_{target} , ε , α , and β , the minimum required depth is uniquely determined. Cabello-Aguilar et al. [2025] propose dynamic LOD calibration based on sequencing depth; the present framework provides the theoretical expression that such calibration should implement.

6.2 Clinical Scenario Benchmarking

Three canonical clinical scenarios are evaluated in Figure 4, using parameters from published validation studies. **Early detection** (target $f \sim 0.01\text{--}0.1\%$): UMI error-corrected NGS at $10,000\times$ meets theoretical requirements, consistent with Frydendahl et al. [2023] demonstrating 81% sensitivity at 0.01% VAF. **MRD post-surgery** (target $f \sim 0.001\text{--}0.01\%$): duplex WGS at $120,000\times$ meets theoretical requirements, consistent with Northcott et al. [2024] achieving $\text{LOD}_{95} = 3.45$ PPM and Kallio et al. [2024] demonstrating ctDNA tracking to 0.0004% in ovarian cancer. **Treatment response** ($\Delta\text{AF} > 0.5\%$): standard large-panel NGS at $1,000\times$ meets theoretical requirements, consistent with Yu et al. [2022] showing $\geq 97\%$ sensitivity at $\geq 0.5\%$ VAF.

All three current clinical practice approaches are theoretically justified—the first formal confirmation that existing depth standards are adequate for their stated clinical applications.

6.3 Regulatory Implications

FDA analytical validation guidance for NGS-based IVDs [US Food and Drug Administration, 2023] requires specification of analytical sensitivity as limit of detection. The present framework provides the theoretical basis for LOD specification: regulatory submissions can present both empirical LOD (from dilution experiments) and theoretical LOD from Equation (7), demonstrating that empirical performance approaches or achieves theoretical limits. This is directly applicable to MHRA SaMD technical file requirements for

Figure 4 — Clinical Scenario Benchmarking: Theoretical Requirements vs Current Practice
Derived from Neyman-Pearson framework validated against US and EU published assay data (2018-2025)

Clinical Scenario	Target f	Recommended Platform	Theoretical N*	f_min at Practice Depth	Practice Depth	Meets Requirement?	Clinical Application	Evidence Base
Early Cancer Detection	0.01-0.5%	UMI NGS ($\epsilon=10^{-5}$)	62x	2.89e-04	10,000x	✓ Yes	Lung, breast, CRC screening	Frydendahl 2023 (Denmark, EU)
MRD Post-Surgery	0.001-0.01%	Duplex WGS ($\epsilon=7.7\times10^{-7}$)	476x	2.42e-05	120,000x	✓ Yes	CRC, ovarian, NSCLC	Northcott 2024 (Personalis, US)
Treatment Response	$\Delta AF > 0.5\%$	Large-panel NGS ($\epsilon=10^{-3}$)	247x	4.51e-03	1,000x	✓ Yes	EGFR/BRAF monitoring	Yu 2022 (BMS, US)

N* = minimum required depth from f_min formula ($\alpha=0.05$, power=0.80). Current practice depths substantially exceed N* — reflecting multi-locus panel integration not captured by the single-locus model. All practice depths empirically validated in cited sources. Vitacore Research | awalkr@vitacore-research.org

Figure 4: **Clinical scenario benchmarking: theoretical requirements vs current practice.** Three canonical clinical applications evaluated using Equation (9), with parameters from published validation studies. Early detection (UMI NGS, Frydendahl et al. 2023); MRD post-surgery (Duplex WGS, Northcott et al. 2024, Kallio et al. 2024); Treatment response (standard NGS, Yu et al. 2022). All three current practice depths meet theoretical minimum requirements, providing the first formal confirmation that existing clinical depth standards are adequate for their stated applications.

ctDNA assays as Software as a Medical Device.

6.4 NHS Genomics Medicine Service Standardisation

The NHS GMS is deploying ctDNA testing across NHS trusts using heterogeneous platforms. The absence of a common performance specification framework—identified by Dao et al. [2023] and Pantel and Alix-Panabières [2024] as a critical barrier—means that assay comparisons are currently made empirically, laboratory against laboratory, without an absolute reference. The ratio $LOD_{\text{empirical}}/f_{\min}(N, \epsilon)$ provides exactly that reference: a platform-independent efficiency metric. Mandating reporting of this ratio as part of NHS ctDNA assay commissioning specifications would provide standardisation currently not achievable by empirical comparison alone.

7 Conclusions and Future Work

This paper has provided the first formal Neyman-Pearson signal detection theory framework for ctDNA detection. Extending the Binomial detection probability approach of Liu et al. [2021] to formal optimality, we derived the uniformly most powerful detector and

the minimum detectable tumour fraction in closed form:

$$f_{\min} \approx (z_{1-\alpha} + z_{1-\beta}) \sqrt{\varepsilon/N}$$

Shannon channel capacity analysis [Shannon, 1948] identifies $C = 1 - H(\varepsilon)$ bits per read as the per-read information ceiling, showing all modern platforms operate within 2% of this ceiling. The optimal sequencing depth $N^* = (z_{1-\alpha} + z_{1-\beta})^2 \varepsilon / f_{\text{target}}^2$ defines the point beyond which additional depth provides diminishing clinical returns.

Empirical validation against $n = 61$ data points from five peer-reviewed published studies [Yu et al., 2022, Stetson et al., 2018, Frydendahl et al., 2023, Northcott et al., 2024, Cheng et al., 2025] yields 84.3% agreement for single-locus assays. Systematic discordance for multi-locus tumour-informed assays is consistent with model assumptions and identifies the multi-locus extension as the primary future research direction.

Future work:

1. Derivation of the joint multi-locus detection limit as a function of panel size L , accounting for correlated mutations and shared tumour-of-origin signal [Northcott et al., 2024]
2. Formal treatment of clonal haematopoiesis as a structured interference source distinct from the Binomial error model [Semenkovich et al., 2023]
3. Empirical validation of multi-locus $f_{\min}^{(L)}$ predictions against published tumour-informed assay sensitivity data across panel sizes
4. Integration of the formal LOD expression into MHRA and FDA analytical validation frameworks [US Food and Drug Administration, 2023]
5. Development of an open-source Vitacore computational tool implementing the formal detection limits for clinical laboratory use

Data Availability

No primary data were generated or analysed in this study. All numerical values used for empirical validation were extracted from published peer-reviewed analytical

validation studies cited in the manuscript. All source papers are publicly available via PubMed Central and PLOS ONE. Analysis code (Python) is available at <https://vitacore-research.org/code/ctdna-detection-limits> upon publication.

Ethics Statement

This study is purely theoretical and mathematical. No human subjects data, animal data, or patient samples were used at any stage. No ethics approval was required. No human data were accessed: validation was performed using published summary statistics from peer-reviewed papers only.

Competing Interests

The author declares no competing interests.

Funding

This work received no external funding and was conducted as part of Vitacore Research's independent open research programme.

Acknowledgements

The author thanks the research teams at Bristol Myers Squibb, Inivata, Aarhus University, Personalis, and the Broad Institute for making their analytical validation data publicly available, enabling this theoretical framework to be empirically tested.

References

Simon Cabello-Aguilar, Julie A. Vendrell, and Jérôme Solassol. Real-world technical hurdles of ctDNA NGS analysis: Lessons from clinical implementation. *Diseases*, 13(10):312, 2025. doi: 10.3390/diseases13100312.

Alexandre P. Cheng et al. Error-corrected flow-based sequencing at whole-genome scale and its application to circulating cell-free DNA profiling. *Nature Methods*, 2025. doi: 10.1038/s41592-024-02576-2.

Jonathan Dao, Neil Finlayson, Dana Bagnato, Anat Gilad-Opher, Alexandra Mayer, Elisa Sturtevant, and Nir Peled. Using cfDNA and ctDNA as oncologic markers: A path to clinical validation. *International Journal of Molecular Sciences*, 24(11):9368, 2023. doi: 10.3390/ijms24119368.

Amanda Frydendahl, Camilla Kronborg, Boe Sandahl Sørensen, Boe Steen Bech, Kell Østerlind, and Niels Pallisgaard. Error-corrected deep targeted sequencing of circulating cell-free DNA from colorectal cancer patients for sensitive detection of circulating tumor DNA. *International Journal of Molecular Sciences*, 24(4):3888, 2023. doi: 10.3390/ijms24043888.

Heidi Kallio et al. Sensitive circulating tumor DNA-based residual disease detection in epithelial ovarian cancer. *Life Science Alliance*, 2024.

Steven M. Kay. *Fundamentals of Statistical Signal Processing, Volume II: Detection Theory*. Prentice Hall, Upper Saddle River, NJ, 1998. ISBN 978-0135041352.

Zhe Liu, Weijin Qiu, Shujin Fu, Xia Zhao, Jun Xia, Chunyu Geng, Youqian Yu, Ziling Li, Mingzhu Li, Hui Jiang, and Fang Chen. Influence of sequencing depth on the fidelity and sensitivity of 1%–5% low-frequency mutation detection and recommendation for standardization of sequencing depth. *bioRxiv*, 2021. doi: 10.1101/2021.12.26.474219.

Iver Nordentoft, Sia Viborg Lindskrog, Karin Birkenkamp-Demtröder, Santiago Gonzalez, Maja Kuzman, Jurica Levatic, Dunja Glavas, Ryan Ptashkin, James Smadbeck, Danielle Afterman, Tomer Lauterman, Yarin Cohen, Zohar Donenhirsh, Iman Tavassoly, Ury Alon, Amanda Frydendahl, Mads Heilskov Rasmussen, Claus Lindbjerg Andersen, Philippe Lamy, Michael Knudsen, Paz Polak, Asaf Zviran, Boris Oklander, Mads Agerbaek, Joergen Bjerggaard Jensen, and Lars Dyrskjot. Whole-genome mutational analysis for tumor-informed detection of circulating tumor DNA in patients with urothelial carcinoma. *European Urology*, 86(4):301–311, 2024. doi: 10.1016/j.eururo.2024.05.014.

- 339 Josette M. Northcott, Gabor Bartha, Jason Harris, Conan Li, Fabio C. P. Navarro,
340 Rachel Marty Pyke, Manqing Hong, Qi Zhang, Shuyuan Ma, Tina X. Chen, Janet Lai,
341 Nitin Udar, Juan-Sebastian Saldivar, Erin Ayash, Joshua Anderson, Jiang Li, Tiange
342 Cui, Tu Le, Ruthie Chow, Randy Jerel Velasco, Rose Santiago, Robert C. Bruce, Lau-
343 rie J. Goodman, Yi Chen, Dan Norton, Richard O. Chen, and John M. Lyle. Analytical
344 validation of NeXT Personal, an ultra-sensitive personalized circulating tumor DNA as-
345 say. *Oncotarget*, 15:200–218, 2024. doi: 10.18632/oncotarget.28565.
- 346 Klaus Pantel and Catherine Alix-Panabières. Minimal residual disease as a target for
347 liquid biopsy in patients with solid tumours. *Nature Reviews Clinical Oncology*, 2024.
348 doi: 10.1038/s41571-024-00967-y.
- 349 Yoshiharu Sato. Liquid biopsy for minimal residual disease and monitoring in early-stage
350 non-small cell lung cancer. *Exploration of Medicine*, 2025.
- 351 Thomas D. Schneider. Information content of individual genetic sequences. *Journal of*
352 *Theoretical Biology*, 189(4):427–441, 1997. doi: 10.1006/jtbi.1997.0540.
- 353 Nicole P. Semenkovich et al. Genomic approaches to cancer and minimal residual disease
354 detection using circulating tumor DNA. *Journal for Immunotherapy of Cancer*, 2023.
- 355 Claude E. Shannon. A mathematical theory of communication. *Bell System Technical*
356 *Journal*, 27(3):379–423, 1948. doi: 10.1002/j.1538-7305.1948.tb01338.x.
- 357 Sarah Slater et al. Tissue-free liquid biopsies for minimal residual disease detection in
358 early colorectal cancer: TRACC part B. *Clinical Cancer Research*, 2024.
- 359 Daniel Stetson, Ahad Ahmed, Xiaowei Xu, Brian R. B. Nuttall, Thomas J. Lubin-
360 ski, James H. Johnson, J. Carl Barrett, Stephen C. Emery, Robert W. Veltri, and
361 Devashish L. Bhatt. Analytical validation of a next generation sequencing liquid biopsy
362 assay for high sensitivity broad molecular profiling. *PLOS ONE*, 13(3):e0193802, 2018.
363 doi: 10.1371/journal.pone.0193802.
- 364 Roberto Tébar-Martínez et al. Strategies for improving detection of circulating tumor
365 DNA using next generation sequencing. *Cancer Treatment Reviews*, 2023.

US Food and Drug Administration. Considerations for the design, development, and analytical validation of next generation sequencing-based in vitro diagnostics. Technical report, US Food and Drug Administration, 2023.

Jonathan C. M. Wan, Katrin Heider, Davina Gale, Sean Murphy, Edward Fisher, Florent Moulere, Andrea Ruiz-Valdepenas, Elizabeth H. Moore, James Morris, Dineika Chandrananda, and Nitzan Rosenfeld. ctDNA monitoring using patient-specific sequencing and integration of variant reads. *Science Translational Medicine*, 12(553):eaaz8084, 2020. doi: 10.1126/scitranslmed.aaz8084.

Lizhi Yu, Gonzalo Lopez, John Rassa, Yixin Wang, Tara Basavanhally, Andrew Browne, Chang-Pin Huang, Lauren Dorsey, Jin Jen, and Sarah Hersey. Direct comparison of circulating tumor DNA sequencing assays with targeted large gene panels. *PLOS ONE*, 17(4):e0266889, 2022. doi: 10.1371/journal.pone.0266889.

Figure 1 — Neyman-Pearson ctDNA Detection: H_0 vs H_1 Distributions
Parameters matched to published assay conditions




Figure 2 — Theoretical Detection Limit vs Sequencing Depth
Empirical validation: n=61 data points, 5 institutions (US and EU), 2018-2025




Figure 3 — Shannon Information-Theoretic Analysis of ctDNA Sequencing Platforms





Figure 4 — Clinical Scenario Benchmarking: Theoretical Requirements vs Current Practice
Derived from Neyman-Pearson framework validated against US and EU published assay data (2018-2025)

Clinical Scenario	Target f	Recommended Platform	Theoretical N*	f_min at Practice Depth	Practice Depth	Meets Requirement?	Clinical Application	Evidence Base
Early Cancer Detection	0.01-0.5%	UMI NGS (e=10 ⁻⁵)	62x	2.89e-04	10,000x	Yes	Lung, breast, CRC screening	Frydendahl 2023 (Aarhus, Denmark)
MRD Post-Surgery	0.001-0.01%	Duplex WGS (e=7.7x10 ⁻⁷)	476x	2.42e-05	120,000x	Yes	CRC, ovarian, NSCLC	Northcott 2024 (Personalis, US)
Treatment Response	DeltaAF >0.5%	Large-panel NGS (e=10 ⁻³)	247x	4.51e-03	1,000x	Yes	EGFR/BRAF monitoring	Yu 2022 (BMS, US)

N* = minimum required depth from f_min formula (a=0.05, power=0.80). Current practice depths substantially exceed N* -- reflecting multi-locus panel integration not captured by the single-locus model. All practice depths empirically validated in cited sources.
Vitacore Research | awalkr@vitacore-research.org