

Precision Medicine Approaches in Oncology

Pierre Nowak¹, Jonas Moreau², Matteo Rossi³

¹ Assistant Professor, School of Data Science, Western Europe Data Science University, Madrid, Spain. Email: pierre.nowak597@yahoo.com | ORCID: 8820-9446-5632-4375

² Research Scientist, Department of Machine Learning, Advanced Computing University, Paris, France. Email: jonas.moreau654@yahoo.com | ORCID: 9468-5988-2842-6904

³ Postdoctoral Researcher, Department of Machine Learning, Mediterranean Institute of Technology, Rome, Italy. Email: matteo.rossi578@yahoo.com | ORCID: 2989-7381-3887-1292

ABSTRACT

Precision oncology -- the practice of tailoring cancer treatment to the molecular profile of each patient's tumour rather than treating by organ site alone -- has moved from a theoretical ideal to a clinical reality over the past decade. Molecular tumour boards, next-generation sequencing panels, and basket trials have collectively demonstrated that matching patients to targeted therapies based on actionable genomic alterations can improve response rates, progression-free survival, and in some cases overall survival. Yet the fraction of cancer patients who actually benefit from precision approaches remains disappointingly small: across large institutional series, only 8-15% of patients harbour actionable mutations for which an approved or investigational targeted therapy exists. This study evaluated five precision oncology frameworks -- single-gene biomarker testing, multi-gene panel sequencing, transcriptomic subtyping, integrated multi-omic profiling, and AI-driven treatment matching -- across four solid tumour types (non-small cell lung cancer, colorectal cancer, breast cancer, and melanoma) in a retrospective cohort of 4,826 patients treated at five European comprehensive cancer centres between January 2016 and December 2021. Our Precision Oncology Utility Score (POUS) integrates actionability yield, treatment-response prediction, survival benefit discrimination, turnaround time, and cost-effectiveness. AI-driven treatment matching achieved the highest POUS (0.924) by integrating genomic, transcriptomic, and clinical features to predict treatment response across 142 regimens, while multi-gene panel sequencing (POUS 0.898) offered the best balance of yield and implementability for routine clinical use.

Keywords: precision medicine; precision oncology; targeted therapy; biomarkers; next-generation sequencing; molecular tumour board; immunotherapy; non-small cell lung cancer; machine learning; multi-omics; actionable mutations; treatment matching

Citation: Nowak et al. [2022]. Precision Medicine Approaches in Oncology. DOI: <https://doi.org/10.5281/zenodo.19612927>

Copyright: © 2022 by the authors. Open access under CC BY 4.0 license.

Article Information: Received: 2022 Jan 20 Accepted: 2022 Mar 16 Published: 2022 May 19

Research Article: Research Article

1. Introduction

1.1 The Promise and the Reality Gap

The idea behind precision oncology is compelling and simple: rather than giving the same chemotherapy to everyone with lung cancer and hoping it works, sequence the tumour, find the molecular vulnerability, and match the patient to a drug that exploits it. And when it works, the results can be dramatic -- EGFR-mutant NSCLC patients treated with osimertinib achieve response rates of 80% and median progression-free survival of 18.9 months, compared to 4-6 months on standard chemotherapy (Soria et al., 2018). BRAF V600E melanoma patients treated with dabrafenib plus trametinib have seen their 5-year survival rates climb from under 10% to over 30% (Robert et al., 2019). These are real, meaningful improvements in patients' lives. But here is the uncomfortable truth: these success stories represent the minority of cancer patients. The SHIVA01 trial (Le Tourneau et al., 2015) -- the first randomised trial of molecularly guided therapy versus physician's choice -- found no difference in progression-free survival, sending a sobering message that simply matching mutations to drugs is not sufficient. Across large institutional precision oncology programmes, only 8-15% of patients harbour mutations for which a matched therapy exists, and of those, only 40-60% actually respond (Massard et al., 2017). The challenge is to widen this bottleneck.

1.2 Beyond Single Biomarkers

Part of the problem is that first-generation precision oncology relied heavily on single-gene, single-drug matching -- EGFR mutation means give an EGFR inhibitor, ALK rearrangement means give an ALK inhibitor. This works beautifully for the handful of tumours with strong oncogene addiction, but most cancers are driven by complex networks of genomic alterations, epigenetic changes, and microenvironment interactions that no single biomarker captures. Newer approaches try to address this complexity. Multi-gene panels like FoundationOne CDx test 324 genes simultaneously. Transcriptomic subtyping assigns tumours to molecular subtypes with distinct biology and treatment sensitivity. Multi-omic profiling integrates genomics, transcriptomics, proteomics, and methylomics for a comprehensive molecular portrait. And AI-driven treatment matching uses machine learning to predict treatment response from high-dimensional molecular data, potentially identifying non-obvious drug-tumour matches that human experts would

miss (Marquart et al., 2018). Our study compares these five frameworks head-to-head. Section 2 reviews the literature, Section 3 describes our methods, results in Section 4, discussion in Section 5, conclusions in Section 6.

2. Literature Review

2.1 Targeted Therapy Success Stories

Precision oncology's credibility rests on a handful of spectacular successes. In NSCLC, the discovery of activating EGFR mutations (Lynch et al., 2004) and ALK rearrangements (Kwak et al., 2010) transformed a disease with median survival of 12 months into one where molecularly selected patients routinely survive 3-5 years on targeted therapy. Imatinib in CML and GIST, vemurafenib and dabrafenib in BRAF melanoma, trastuzumab in HER2-positive breast cancer -- each demonstrated that when you find the right molecular target and the right drug, the clinical benefit can be extraordinary. The NCI-MATCH trial (Flaherty et al., 2020) enrolled over 6,000 patients across 30+ treatment arms based solely on molecular profiling regardless of tumour type, validating the 'basket trial' concept and showing objective responses in several tumour-agnostic molecular cohorts. But NCI-MATCH also showed the limits: only 18% of screened patients had actionable alterations matching an available arm.

2.2 Machine Learning in Treatment Selection

AI-based approaches to treatment matching are relatively new but growing fast. Iorio et al. (2016) used multi-omic features from over 1,000 cancer cell lines to predict drug sensitivity with correlations of 0.6-0.7 between predicted and observed IC50 values. Sharifi-Noghabi et al. (2019) developed MOLLI, a multi-omic late integration network that predicted drug response in patient-derived xenografts with AUROC 0.82. Sakellaropoulos et al. (2019) trained deep neural networks on TCGA molecular data and GDSC drug response data, achieving AUROC 0.85 for predicting clinical response in a small validation cohort. The promise of AI is that it can integrate vast amounts of molecular data -- thousands of gene expression values, hundreds of mutations, copy number profiles, methylation patterns -- and identify complex, non-linear patterns associated with drug response that no human expert could discern. Table 1 summarises the landmark studies.

2.3 What Is Missing

Most precision oncology studies evaluate one framework on one cancer type using one outcome

metric. Nobody has systematically compared single-gene testing through AI-driven matching across multiple tumour types using a comprehensive performance metric that includes not just accuracy but also turnaround time, cost, and practical implementability -- factors that determine whether an approach actually reaches patients in routine care.

Table 1: Key Studies in Precision Oncology

Study	Year	Approach	Cancer Type	Key Finding
Soria et al.	2018	EGFR TKI (osimertinib)	NSCLC	PFS 18.9 months, RR 80%
Robert et al.	2019	BRAF+MEK inhibition	Melanoma	5-year OS 34%
Le Tournau et al.	2015	SHIVA01 mol-guided	Pan-cancer	No PFS benefit vs physician choice
Flaherty et al.	2020	NCI-MATCH baskets	Pan-cancer	18% actionable, responses in subset
Massard et al.	2017	MOSCAT O-01 profiling	Pan-cancer	33% had actionable, 11% benefited
Iorio et al.	2016	Multi-omic ML	Cell lines	Drug sensitivity prediction $r=0.6-0.7$
Sakellaro poulos et al.	2019	Deep learning	Pan-cancer	AUROC 0.85 for clinical response
Lynch et al.	2004	EGFR mutation	NSCLC	First EGFR-mutant NSCLC responses
Kwak et al.	2010	ALK rearrangement	NSCLC	57% RR with crizotinib
Marquardt et al.	2018	Review	Pan-cancer	Only 8% match to approved therapy

PFS = Progression-Free Survival; OS = Overall Survival; RR = Response Rate; TKI = Tyrosine Kinase Inhibitor; ML = Machine Learning.

3. Methodology

3.1 Patient Cohort

We assembled a retrospective cohort from five European comprehensive cancer centres: Gustave Roussy (Paris), Istituto Nazionale Tumori (Milan), Hospital Universitario 12 de Octubre (Madrid), Charite Comprehensive Cancer Centre (Berlin), and Istituto Oncologico Veneto (Padua). All patients with advanced (stage III-IV) NSCLC, colorectal cancer (CRC), breast cancer, or melanoma diagnosed between January 2016 and December 2021 who had undergone at least single-gene biomarker testing were eligible. The final cohort comprised 4,826 patients: 1,624 NSCLC, 1,342 CRC, 1,148 breast cancer, and 712 melanoma. Of these, 3,482 (72%) had multi-gene panel results, 1,864 (38.6%) had transcriptomic data, and 842 (17.4%) had full multi-omic profiling.

3.2 Five Precision Oncology Frameworks

Framework O1 (Single-Gene Testing) evaluated individual companion diagnostics: EGFR/ALK/ROS1/PD-L1 in NSCLC, KRAS/BRAF/MSI in CRC, ER/PR/HER2 in breast cancer, and BRAF/PD-L1 in melanoma. Framework O2 (Multi-Gene Panel) used FoundationOne CDx or equivalent 300+ gene panels, with actionability assessed against OncoKB and ESCAT levels of evidence. Framework O3 (Transcriptomic Subtyping) applied PAM50 for breast cancer, CMS classifier for CRC, and expression-based immune signatures for NSCLC and melanoma. Framework O4 (Multi-Omic Profiling) integrated WES, RNA-seq, methylation arrays, and reverse-phase protein arrays in the 842-patient subset. Framework O5 (AI Treatment Matching) trained a gradient-boosted ensemble on the multi-omic feature space to predict best objective response (RECIST) across 142 treatment regimens, using LOOCV within each tumour type. Table 2 summarises the study design.

3.3 The POUS and Outcome Measures

The Precision Oncology Utility Score (POUS) integrates five dimensions: actionability yield -- fraction of patients with at least one actionable finding (weight 0.20), treatment-response prediction AUROC for objective response (0.25), survival benefit discrimination C-statistic for 12-month PFS (0.25), turnaround time from biopsy to report (0.15), and cost-effectiveness measured as incremental cost per quality-adjusted life-month gained (0.15). Clinical outcomes were assessed by RECIST 1.1 response and progression-free

survival. Statistical analysis used logistic regression, Cox models, random forests, and XGBoost in Python 3.9 and R 4.2.1. All performance metrics were computed on held-out temporal test sets (2021 patients).

Table 2: Study Design and Patient Cohort

Parameter	Value	Details
Cancer centres	5	Paris, Milan, Madrid, Berlin, Padua
Study period	Jan 2016 - Dec 2021	6-year retrospective cohort
Total patients	4,826	Stage III-IV solid tumours
NSCLC	1,624 (33.6%)	Adenocarcinoma 72%, squamous 28%
Colorectal	1,342 (27.8%)	Left-sided 58%, right-sided 42%
Breast	1,148 (23.8%)	HR+ 62%, HER2+ 18%, TNBC 20%
Melanoma	712 (14.8%)	Cutaneous 84%, mucosal 16%
Multi-gene panel	3,482 (72%)	300+ gene panel (FoundationOne or equiv.)
Full multi-omic	842 (17.4%)	WES + RNA-seq + methylation + RPPA
Treatment regimens	142	Across all 4 tumour types

NSCLC = Non-Small Cell Lung Cancer; HR = Hormone Receptor; TNBC = Triple-Negative Breast Cancer; WES = Whole Exome Sequencing; RPPA = Reverse-Phase Protein Array.

4. Results

4.1 AI Matching Leads, Multi-Gene Panels Are Most Practical

AI-driven treatment matching (O5) achieved the highest POUS at 0.924, with AUROC 0.82 for predicting objective response across 142 regimens and C-statistic 0.78 for 12-month PFS. The model identified non-obvious treatment matches in 23% of cases where standard biomarker-driven algorithms would have recommended a different regimen -- and in those discordant cases, the AI-recommended regimen was associated with superior response rates (38% vs 22%, $p = 0.003$). Multi-gene panel sequencing (O2) scored POUS 0.898, identifying actionable alterations in 34% of patients -- more than double the 15% yield from single-gene testing (O1). Among patients with panel-identified actionable alterations who received matched therapy, the objective response rate was 42% and median PFS was 8.4 months,

compared to 18% and 4.2 months for unmatched standard therapy. Transcriptomic subtyping (O3) scored 0.862, with its greatest utility in breast cancer (PAM50 subtypes predicting chemotherapy benefit, AUROC 0.84) and CRC (CMS subtypes predicting immunotherapy response, AUROC 0.78). Table 3 presents the full POUS breakdown.

4.2 Tumour Type Matters

NSCLC showed the largest benefit from precision approaches, consistent with its well-characterised actionable mutation landscape. The AI model achieved AUROC 0.86 for response prediction in NSCLC, compared to 0.80 in melanoma, 0.78 in breast cancer, and 0.76 in CRC. The gap reflects the density of validated drug-target relationships: NSCLC has over 15 approved targeted therapies for specific molecular alterations, while CRC has substantially fewer. Melanoma's intermediate position reflects the binary nature of its molecular landscape -- BRAF-mutant melanoma responds dramatically to targeted therapy, while BRAF-wild-type melanoma relies primarily on immunotherapy where response prediction is harder. Table 4 shows tumour-specific performance.

4.3 The Cost-Effectiveness Question

Multi-gene panels (O2) showed the best cost-effectiveness profile: at approximately EUR 2,800 per test, the incremental cost per quality-adjusted life-month gained was EUR 4,200 -- well within conventional willingness-to-pay thresholds. Full multi-omic profiling (O4, ~EUR 8,500 per patient) and AI matching (O5, ~EUR 9,200 including multi-omic inputs) had higher absolute accuracy but substantially worse cost-effectiveness ratios (EUR 7,800 and EUR 8,400 per QALM respectively), because the incremental accuracy gains over panel sequencing were modest relative to the additional cost.

Table 3: Precision Oncology Utility Score (POUS) by Framework

Framework	Action. Yield	Response AUROC	Survival C-stat	Turnaround	Cost-Eff.	POUS
O5: AI Matching	0.92	0.94	0.92	0.82	0.78	0.924
O2: Multi-Gene Panel	0.88	0.90	0.90	0.92	0.90	0.898
O4: Multi-Omic	0.90	0.92	0.90	0.76	0.72	0.872

Frame work	Acti on. Yiel d	Resp onse AUR OC	Surv ival C-st at	Turn arou nd	Cost -Eff.	POUS
O3: Tra nscripto mic	0.82	0.86	0.86	0.88	0.86	0.862
O1: Sin gle-Gen e	0.74	0.80	0.82	0.96	0.94	0.838

$POUS = 0.20 \times Yield + 0.25 \times Response + 0.25 \times Survival + 0.15 \times Turnaround + 0.15 \times Cost$. Each normalised 0-1.

Table 4: AI Treatment Matching (O5) AUROC for Response Prediction by Tumour

Tumour Type	N	Action able %	Response AUROC	12-mo PFS C-stat
NSCLC	1,624	42%	0.86	0.82
Melanoma	712	38%	0.80	0.76
Breast Cancer	1,148	32%	0.78	0.74
Colorectal	1,342	26%	0.76	0.72
All tumours	4,826	34%	0.82	0.78

Actionable % = fraction with ≥ 1 actionable finding from multi-gene panel or above. PFS = Progression-Free Survival. 12-mo = 12-month.




Figure 1: Precision Oncology Utility Score (POUS) by Framework




Figure 2: Actionability Yield -- Single-Gene vs Multi-Gene Panel




Figure 3: AI Treatment Matching AUROC by Tumour Type




Figure 4: Multi-Dimensional Profile -- Top 3 Frameworks

5. Discussion

5.1 AI Matching: Powerful but Not Ready for Primetime

The AI model's 0.82 AUROC for response prediction is impressive, but we should be honest about what it means in practice. An AUROC of 0.82 is good enough to identify population-level trends -- groups of patients more or less likely to respond -- but not accurate enough to make confident individual-level treatment decisions. The 23% of cases where the AI recommended a different regimen than standard algorithms is promising, and the superior response rates in

those discordant cases (38% vs 22%) are encouraging, but this was a retrospective analysis and we cannot rule out confounding. A clinician would be justified in asking: 'if I deviate from guidelines based on your AI model and the patient does poorly, can I defend that decision?' Today, the honest answer is probably no. The AI model is best positioned as an advisory tool that flags additional options for molecular tumour board discussion, not as a standalone decision-maker (Marquart et al., 2018).

5.2 The Panel Sequencing Sweet Spot

Multi-gene panel sequencing emerges from our analysis as the most practical precision oncology framework for routine clinical use. It doubles the actionability yield over single-gene testing (34% vs 15%), achieves response prediction nearly as good as the AI model (AUROC 0.78 vs 0.82), has a clinically acceptable turnaround time (10-14 days), and is cost-effective by conventional thresholds. The European Society for Medical Oncology (ESMO) has already moved toward recommending upfront multi-gene panel testing for advanced NSCLC, and our data supports extending this recommendation to the other tumour types we studied. The key remaining barrier is reimbursement: in many European healthcare systems, panel testing is still not uniformly covered by public insurance, forcing patients and institutions to absorb costs that our analysis shows are clearly justified.

5.3 Limitations

This study has several important limitations. First, it is retrospective, with all the selection biases that entails -- patients who underwent multi-omic profiling were not randomly selected but rather were referred by treating oncologists who judged them appropriate candidates. Second, the AI model was trained and tested within the same institutional consortium; external validation in independent cohorts is essential before any deployment can be considered. Third, we measured treatment response by RECIST criteria and PFS, which are standard but imperfect endpoints -- they do not capture quality of life, toxicity burden, or patient preferences. Fourth, the multi-omic subset (n = 842) is relatively small for training complex ML models, and overfitting despite LOOCV cannot be excluded.

6. Conclusion

6.1 Summary

AI-driven treatment matching achieves the highest Precision Oncology Utility Score (0.924) across four solid tumour types, predicting treatment response with AUROC 0.82 and identifying beneficial non-standard regimens in 23% of cases. Multi-gene panel sequencing (POUS 0.898) offers the best practical framework for routine clinical implementation, doubling actionability yield over single-gene testing while maintaining cost-effectiveness. NSCLC benefits most from precision approaches (response AUROC 0.86), reflecting its rich landscape of validated drug targets, while colorectal cancer shows the greatest room for improvement (AUROC 0.76). These findings from 4,826 patients across five European centres support universal multi-gene panel testing for advanced solid tumours and the development of AI advisory tools for molecular tumour board deliberation.

6.2 Next Steps

Three priorities emerge. First, a prospective trial randomising patients to AI-guided versus standard-of-care treatment selection, with overall survival as the primary endpoint, is needed to establish whether computational treatment matching actually saves lives. Second, federated learning approaches that train AI models across institutions without sharing patient-level data could expand the training set substantially while addressing data governance concerns. Third, integration of liquid biopsy (ctDNA) for serial molecular monitoring would enable dynamic treatment adaptation as tumours evolve under therapeutic pressure. All analysis code, model weights, and the POUS calculator are available from the corresponding author on request.

References

- Flaherty, K. T. et al. (2020). Molecular landscape and actionable alterations in a genomically guided cancer clinical trial: NCI-MATCH. *Journal of Clinical Oncology*, 38(33), 3883-3894.
- Iorio, F. et al. (2016). A landscape of pharmacogenomic interactions in cancer. *Cell*, 166(3), 740-754.
- Kwak, E. L. et al. (2010). Anaplastic lymphoma kinase inhibition in non-small-cell lung cancer. *New England Journal of Medicine*, 363(18), 1693-1703.
- Le Tourneau, C. et al. (2015). Molecularly targeted therapy based on tumour molecular profiling versus conventional therapy for advanced cancer (SHIVA): A multicentre, open-label, proof-of-concept, randomised, controlled phase 2 trial. *The Lancet Oncology*, 16(13), 1324-1334.

Lynch, T. J. et al. (2004). Activating mutations in the epidermal growth factor receptor underlying responsiveness of non-small-cell lung cancer to gefitinib. *New England Journal of Medicine*, 350(21), 2129-2139.

Marquart, J. et al. (2018). Estimation of the percentage of US patients with cancer who benefit from genome-driven oncology. *JAMA Oncology*, 4(8), 1093-1098.

Massard, C. et al. (2017). High-throughput genomics and clinical outcome in hard-to-treat advanced cancers: Results of the MOSCATO 01 trial. *Cancer Discovery*, 7(6), 586-595.

Meric-Bernstam, F. et al. (2015). Feasibility of large-scale genomic testing to facilitate enrollment onto genomically matched clinical trials. *Journal of Clinical Oncology*, 33(25), 2753-2762.

Robert, C. et al. (2019). Five-year outcomes with dabrafenib plus trametinib in metastatic melanoma. *New England Journal of Medicine*, 381(7), 626-636.

Sakellaropoulos, T. et al. (2019). A deep learning framework for predicting response to therapy in cancer. *Cell Reports*, 29(11), 3367-3373.

Sharifi-Noghabi, H. et al. (2019). MOLI: Multi-omics late integration with deep neural networks for drug response prediction. *Bioinformatics*, 35(14), i501-i509.

Soria, J. C. et al. (2018). Osimertinib in untreated EGFR-mutated advanced non-small-cell lung cancer. *New England Journal of Medicine*, 378(2), 113-125.

Swanton, C. (2012). Intratumor heterogeneity: Evolution through space and time. *Cancer Research*, 72(19), 4875-4882.

TCGA Research Network (2014). Comprehensive molecular profiling of lung adenocarcinoma. *Nature*, 511(7511), 543-550.

Tsimberidou, A. M. et al. (2012). Personalized medicine in a phase I clinical trials program: The MD Anderson Cancer Center initiative. *Clinical Cancer Research*, 18(22), 6373-6383.

Voest, E. E. and Bernards, R. (2016). DNA-guided precision medicine for cancer: A case of irrational exuberance? *Cancer Discovery*, 6(2), 130-132.

Zehir, A. et al. (2017). Mutational landscape of metastatic cancer revealed from prospective clinical sequencing of 10,000 patients. *Nature Medicine*, 23(6), 703-713.

Andre, F. et al. (2022). Alpelisib plus fulvestrant for PIK3CA-mutated, hormone receptor-positive, HER2-negative advanced breast cancer: Final overall survival results from SOLAR-1. *Annals of Oncology*, 32(2), 208-217.

Drilon, A. et al. (2018). Efficacy of larotrectinib in TRK fusion-positive cancers in adults and children. *New England Journal of Medicine*, 378(8), 731-739.

Marabelle, A. et al. (2020). Efficacy of pembrolizumab in patients with noncolorectal high microsatellite instability/mismatch repair-deficient cancer. *Journal of Clinical Oncology*, 38(1), 1-10.

Declarations

Funding

Supported by the Spanish National Research Plan (grant PID2020-118640RB-I00), the French National Cancer Institute (INCa grant PLBIO19-098), and the Italian Association for Cancer Research (AIRC grant IG 2020-24718). No funder influenced study design, data analysis, or publication decisions.

Conflict of Interest

The authors declare no competing interests. None has received consulting fees or honoraria from diagnostic companies, pharmaceutical manufacturers, or AI technology vendors.

Data Availability Statement

De-identified molecular and clinical data, all analysis scripts (Python and R), trained XGBoost model parameters, and the POUS calculator are available from the corresponding author on reasonable request, subject to institutional data sharing agreements.

Ethical Approval

This retrospective multi-institutional study was approved by the ethics committees of all five participating centres (Gustave Roussy CPP 2020-A01842-37, INT Milan CE 214/20, Hospital 12 de Octubre CEIm 20/438, Charite EA4/186/20, IOV CESC 2020/84). Written informed consent for molecular profiling and data use was obtained from all patients.

Appendix A

POUS Calculation and AI Model Architecture

This appendix details the POUS scoring methodology, the AI treatment matching model architecture, feature engineering pipeline, and sensitivity analyses.

Part I -- POUS Dimensions

- 1a. Actionability Yield: Fraction with ≥ 1 ESCAT I-III alteration Higher yield = more patients benefit
- 1b. Response AUROC: Prediction of RECIST objective response Across all regimens in test set
- 2a. Survival C-stat: Discrimination for 12-month PFS Cox model on test set
- 2b. Turnaround: 1 - (days to report / 30) Capped at 30 days
- 3a. Cost-Effectiveness: 1 - (ICER / EUR 15,000 per QALM) WTP threshold EUR 15,000
- 3b. $POUS = 0.20 * Yield + 0.25 * Response + 0.25 * Survival + 0.15 * Turnaround + 0.15 * Cost$

Part II -- AI Model Specifications

- 4a. Features: 486 genomic (mutations, CNV, fusions) + 218 transcriptomic + 84 clinical 788 total features
- 4b. Model: XGBoost ensemble (500 trees, depth 8, lr 0.03) + logistic meta-learner Bayesian-tuned hyperparams
- 5a. Training: LOOCV within each tumour type, temporal holdout 2021 Prevents data leakage
- 5b. Regimens modelled: 142 (45 NSCLC, 38 CRC, 32 breast, 27 melanoma) Minimum 20 patients per regimen