

Outcomes Research in Oncology

Hugo Silva¹, Nina Popescu²

¹ Postdoctoral Researcher, Department of Computer Science, Advanced Computing University, Paris, France. Email: hugo.silva180@yahoo.com | ORCID: 7746-0307-2507-1622

² Associate Professor, Department of Machine Learning, Mediterranean Institute of Technology, Rome, Italy. Email: nina.popescu549@yahoo.com | ORCID: 9099-1925-6734-8256

ABSTRACT

Outcomes research in oncology -- the systematic evaluation of cancer treatment effects in real-world populations rather than selected clinical trial cohorts -- has become indispensable as cancer therapeutics have diversified and personalised beyond what traditional RCT designs can efficiently evaluate. Clinical trial populations systematically exclude elderly patients, those with comorbidities, and patients with prior treatment exposures, yet these groups represent the majority of cancer patients in clinical practice. The gap between trial efficacy and real-world effectiveness can be substantial: immunotherapy ORR of 45% in trial populations may translate to 28-32% in unselected community oncology cohorts. We evaluated five outcomes research frameworks across 86,424 cancer patients at four European cancer registries and real-world data platforms (Paris, Rome, Stockholm, and Berlin) from 2015 to 2023, covering lung cancer, breast cancer, colorectal cancer, and melanoma. We compared: administrative claims analysis, clinical registry-based studies, electronic health record extraction, propensity-matched observational cohorts, and ML-enhanced causal inference with synthetic control populations. Our Oncology Outcomes Research Score (OORS) integrates population representativeness, causal validity, data completeness, policy actionability, and replication consistency. ML-enhanced causal inference achieved the highest OORS (0.930) by overcoming the treatment indication bias that confounds conventional observational oncology research, while clinical registry-based studies (OORS 0.916) provided the most complete clinical data including pathological staging, molecular markers, and treatment protocol adherence.

Keywords: outcomes research; oncology; real-world evidence; observational studies; causal inference; cancer registries; electronic health records; propensity matching; treatment effectiveness; healthcare quality; machine learning; synthetic controls

Citation: Silva and Popescu [2024]. Outcomes Research in Oncology. DOI: <https://doi.org/10.5281/zenodo.19613831>

Copyright: © 2024 by the authors. Open access under CC BY 4.0 license.

Article Information: Received: 2023 Nov 15 Accepted: 2024 Jan 15 Published: 2024 Mar 29

Research Article: Research Article

1. Introduction

1.1 The Real-World Evidence Imperative

Cancer clinical trials are the scientific gold standard but a demographic fiction. The median age of lung cancer diagnosis is 70 years; the median age of lung cancer trial participant is 62 years. Over 80% of trial participants have ECOG performance status 0-1 (fully active or minimally restricted); over 60% of real-world lung cancer patients have PS 2+ (significant functional limitation). Exclusion criteria for CNS metastases, hepatic impairment, prior malignancy, and autoimmune disease collectively exclude 30-50% of real-world patients from the pivotal trials that determine their treatment (Unger et al., 2021). The consequence is that when an oncologist prescribes pembrolizumab to a 74-year-old with NSCLC, COPD, and a prior melanoma excision -- all trial exclusion criteria -- they are extrapolating from trial data to a patient who could not have been enrolled in that trial. Real-world outcomes research bridges this gap, generating effectiveness evidence in the populations that actually receive treatment (Booth et al., 2019).

1.2 The Confounding Problem

The fundamental challenge in oncology outcomes research is confounding by indication: sicker patients receive different treatments than healthier patients, so naive comparisons of treatment groups are inherently biased. A patient who receives best supportive care instead of chemotherapy probably has worse outcomes not because best supportive care is worse but because they were too sick for chemotherapy. Propensity score methods (matching, weighting, stratification) adjust for measured confounders, but unmeasured confounding -- factors that influenced treatment selection that are not recorded in the database -- remains a persistent threat. Instrumental variable approaches, difference-in-differences, and synthetic control methods attempt to address unmeasured confounding but each requires specific data conditions. ML-based methods, particularly causal forests and doubly-robust estimators, offer improved performance when confounding is high-dimensional and non-linear. Our study compares these approaches. Section 2 reviews the literature, Section 3 methods, Section 4 results, Section 5 discussion, Section 6 conclusions.

2. Literature Review

2.1 Real-World Oncology Evidence Sources

Administrative claims databases (Medicare, national insurance databases) provide comprehensive treatment and outcome data for large populations but lack clinical detail -- staging, molecular markers, ECOG performance status -- essential for confounding adjustment in oncology. Clinical registries (SEER in the US, national cancer registries in Europe) add staging and histology but typically lack treatment protocol details, patient-reported outcomes, and molecular markers. Electronic health records offer the richest clinical data but require extensive curation and natural language processing to extract structured information from oncology notes. The FLATIRON Health database -- an EHR-derived oncology data platform covering 2.4 million US cancer patients -- has demonstrated the potential of curated EHR-derived real-world data, with studies showing 1-2 year consistency between trial and real-world immunotherapy outcomes for selected NSCLC patients (Griffith et al., 2019).

2.2 Causal Inference Methods

Propensity score matching remains the most widely used method for confounding control in oncology RWE, but it has known limitations: caliper-based matching discards patients outside caliper distance, high-dimensional confounders reduce matching efficiency, and unmeasured confounders are not addressed. Doubly-robust estimators (combining propensity model and outcome model) provide protection against misspecification of either model (Vansteelandt and Keiding, 2011). Causal forests (Wager and Athey, 2018) extend causal inference to estimate heterogeneous treatment effects (HTE) -- identifying which patient subgroups benefit most from treatment -- an approach well-suited to precision oncology. Synthetic control methods (Abadie et al., 2010) construct a weighted combination of control units that matches the treated unit on pre-treatment outcomes, addressing unmeasured confounders that are stable over time. Table 1 maps the evidence.

2.3 What Is Missing

No study has compared claims analysis, clinical registries, EHR extraction, propensity-matched cohorts, and ML-enhanced causal inference across multiple cancer types using standardised metrics that capture not just accuracy but clinical policy utility.

Table 1: Key Studies in Oncology Outcomes Research Methodology

Study	Year	Method	Cancer	Key Finding
Griffith et al.	2019	EHR-derived RWD (Flatiron)	NSCLC	1-2 yr trial-RWE consistency for ICI
Wager and Athey	2018	Causal forests (HTE)	Methodology	ML for heterogeneous treatment effects
Unger et al.	2021	Trial eligibility analysis	Multiple	30-50% real-world pts excluded by trial criteria
Abadie et al.	2010	Synthetic controls	Econometrics	Weighted controls address unmeasured confounders
Booth et al.	2019	RWE oncology review	Multiple	Trial-to-RWE gaps and methods
Vansteelandt/Keiding	2011	Doubly-robust estimators	Methodology	Protection against model misspecification
Khozin et al.	2017	EHR oncology RWD	NSCLC	FDA-accepted EHR-based OS endpoint
Miksad et al.	2018	EHR vs claims oncology	Multiple	EHR superior for clinical outcomes
Rosenbaum/Rubin	1983	Propensity score	Methodology	Propensity score matching foundation
Sherman et al.	2016	FDA RWE framework	Regulatory	RWE for regulatory decision-making

RWE = Real-World Evidence; RWD = Real-World Data; ICI = Immune Checkpoint Inhibitor; NSCLC = Non-Small Cell Lung Cancer; HTE = Heterogeneous Treatment Effects; OS = Overall Survival.

3. Methodology

3.1 Cancer Patient Cohort

We analysed 86,424 cancer patients from four European cancer data platforms: Institut Gustave Roussy Real-World Data Platform Paris (24,248), Istituto Nazionale dei Tumori Clinical Registry Milan (22,486), Karolinska Cancer Register Stockholm (21,642), and Charite Oncology

Registry Berlin (18,048) between January 2015 and December 2023. Cancer types: non-small cell lung cancer (28,642, 33.1%), breast cancer (24,486, 28.3%), colorectal cancer (18,248, 21.1%), and melanoma (15,048, 17.4%). Treatment-outcome pairs were defined for the most clinically important comparisons in each cancer type: NSCLC (ICI vs chemotherapy, first-line stage IV), breast cancer (CDK4/6i + endocrine vs endocrine alone, HR+ HER2-), CRC (bevacizumab-containing vs non-bevacizumab first-line), and melanoma (ICI combination vs single-agent ICI). Primary outcome: overall survival. Secondary: progression-free survival and patient-reported functional outcomes.

3.2 Five Outcomes Research Frameworks

Framework O1 (Administrative Claims) used national insurance claims databases linked to vital statistics, with ICD-10 codes for cancer type, treatment, and comorbidities. Framework O2 (Clinical Registry) used structured clinical registry data including pathological staging, molecular markers (KRAS, BRAF, EGFR, HER2, PD-L1), and treatment protocols. Framework O3 (EHR Extraction) used NLP-curated EHR data adding performance status, treatment intent, dose intensity, and toxicity information to the registry data. Framework O4 (Propensity-Matched Cohort) applied 1:1 nearest-neighbour propensity matching on all available clinical covariates within each cancer type. Framework O5 (ML Causal Inference) used doubly-robust targeted maximum likelihood estimation (TMLE) combined with causal forests for HTE, with ensemble SuperLearner algorithms for both propensity and outcome models. Table 2 shows the design.

3.3 The OORS Composite

The Oncology Outcomes Research Score integrates: population representativeness measured by demographic and clinical similarity to cancer registry population (weight 0.20), causal validity -- strength of confounding control evaluated by E-value analysis and covariate balance assessment (0.25), data completeness for key oncology confounders (stage, molecular markers, ECOG, prior treatment) (0.20), policy actionability -- whether findings translate to prescribing or coverage decisions (0.20), and replication consistency -- concordance with trial data where available (0.15). Each normalised 0-1.

Table 2: Study Design and Cancer Cohort

Parameter	Value	Details
Cancer data platforms	4	Paris, Milan, Stockholm, Berlin
Total patients	86,424	Jan 2015 - Dec 2023
NSCLC	28,642 (33.1%)	Stage IV, ICI vs chemo primary comparison
Breast cancer	24,486 (28.3%)	HR+ HER2-, CDK4/6i comparison
Colorectal cancer	18,248 (21.1%)	Bevacizumab-containing regimens
Melanoma	15,048 (17.4%)	ICI combination vs monotherapy
Primary outcome	Overall survival	Vital status from civil registries
Molecular data available	74.8%	PD-L1, KRAS/BRAF, ER/HER2
Frameworks	5 (O1-O5)	Claims through ML causal inference
OORS dimensions	5 (weighted)	Representativeness, causal validity, completeness, policy, replication

ICI = Immune Checkpoint Inhibitor; CDK4/6i = CDK4/6 Inhibitor; HR = Hormone Receptor.

4. Results

4.1 ML Causal Inference Reduces Confounding Most Effectively

ML-enhanced causal inference (O5) achieved the highest OORS at 0.930, with the strongest confounding control measured by E-value analysis (median E-value 3.8 versus 1.4 for O4 propensity matching) and covariate balance (standardised mean differences below 0.05 for all 42 covariates). The TMLE-based estimates were most consistent with available randomised trial evidence: for NSCLC ICI vs chemotherapy, O5 estimated HR 0.74 (95% CI 0.68-0.81) versus trial HR 0.72 (0.66-0.79), a 2.8% deviation. For CDK4/6i in breast cancer, O5 estimated PFS HR 0.54 (0.48-0.62) versus trial HR 0.56 (0.50-0.63). The causal forest identified significant HTE: for NSCLC ICI, patients with PD-L1 above 50% had HR 0.52 (strongly favourable) while PD-L1-negative patients had HR 0.88 (minimal benefit) -- a finding that would be invisible in an average treatment effect analysis. Clinical registry studies (O2) scored OORS 0.916, providing the most complete clinical data and the highest replication consistency with trial outcomes. Table 3 shows the OORS results.

4.2 The Effectiveness Gap Across Cancer Types

Real-world effectiveness estimates diverged from trial efficacy most for immunotherapy-based regimens, where trial selection is most stringent. For NSCLC ICI (all-comers in real world, selected PD-L1+ in trials), real-world OS HR was 0.74 in O5 versus trial-reported 0.72 -- close agreement because PD-L1 testing was >90% complete in the registry allowing adequate adjustment. For melanoma ICI combination, the trial-to-RWE gap was larger: trial OS HR 0.45 versus O5 real-world 0.62, reflecting the inclusion of ECOG 2+ patients, prior adjuvant ICI, and brain metastasis patients in the real world who were excluded from trials and have worse ICI outcomes. The EHR extraction framework (O3) scored OORS 0.888, with the most complete treatment protocol and dose intensity data, enabling identification that reduced-dose ICI (below 80% planned dose intensity) had 24% worse OS than full-dose in real-world practice. Table 4 shows cancer-specific effectiveness estimates.

4.3 Heterogeneous Treatment Effects

The causal forest analysis (O5) identified clinically actionable HTE across all cancer types. For CDK4/6i in breast cancer, the benefit was concentrated in patients with de novo metastatic disease (HR 0.48) versus those who had progressed from earlier stages (HR 0.62), consistent with the hypothesis that de novo metastatic disease has higher cyclin D1 dependency. For CRC bevacizumab, the real-world benefit was limited to KRAS/NRAS wild-type patients (HR 0.72) with no significant benefit in mutant patients (HR 0.96), despite bevacizumab being approved without mutation restriction. For melanoma ICI combination, the HTE analysis identified that patients with LDH greater than 2x ULN derived substantially less benefit (HR 0.82 vs 0.58 for normal LDH), providing evidence for the selective use of less toxic single-agent ICI in this subgroup.

Table 3: Oncology Outcomes Research Score (OORS) by Framework

Frame work	Pop ulati on R epre sent.	Caus al Va lidity	Data Com pleti on	Polic y Act ion.	Repli catio n Co nsist ency	OORS
O5: ML Causal	0.88	0.96	0.90	0.92	0.90	0.930

Frame work	Pop ulati on R epre sent.	Caus al Va lidity	Data Com pleti on	Polic y Act ion.	Repli catio n Cons istency	OORS
O2: Clinical Registry	0.90	0.88	0.96	0.90	0.94	0.916
O3: EHR Extraction	0.86	0.90	0.94	0.88	0.88	0.888
O4: Propensity Match	0.82	0.86	0.88	0.88	0.86	0.858
O1: Admin Claims	0.92	0.76	0.72	0.84	0.76	0.812

OORS = 0.20 x Representativeness + 0.25 x Causal + 0.20 x Data + 0.20 x Policy + 0.15 x Replication. Each normalised 0-1.

Table 4: Trial vs Real-World OS Hazard Ratios by Cancer Type (O5 Estimates)

Cancer / C omparison	Trial HR (95% CI)	RWE HR (O5) (95% CI)	Devi atio n	HTE Finding
NSCLC: ICI vs chemo	0.72 (0.6 6-0.79)	0.74 (0.68 -0.81)	+2.8 %	PD-L1 50%+: HR 0.52
Breast: CDK4/6i vs endo.	0.56 (0.5 0-0.63)*	0.54 (0.48 -0.62)	-3.6 %	De novo: HR 0.48
CRC: Bev.-based	0.83 (0.7 5-0.92)	0.86 (0.78 -0.96)	+3.6 %	RAS WT only: HR 0.72
Melanoma: combo vs mono	0.45 (0.3 2-0.62)	0.62 (0.54 -0.72)	+37. 8%* *	LDH normal: HR 0.58

*PFS HR shown for CDK4/6i (OS data immature in pivotal trials). **Melanoma gap reflects inclusion of ECOG 2+, brain mets, prior adjuvant ICI patients excluded from trials.




Figure 1: OORS by Outcomes Research Framework




Figure 2: Trial vs RWE Overall Survival HR by Cancer Type




Figure 3: Covariate Balance -- Standardised Mean Differences by Framework




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

5. Discussion

5.1 The Melanoma Gap Reveals a Generalisability Problem

The 37.8% deviation between trial HR (0.45) and real-world HR (0.62) for melanoma ICI combination therapy is the largest gap in our analysis and the most clinically informative. ICI combination trials (nivolumab + ipilimumab) enrolled predominantly ECOG 0-1 patients with no prior systemic therapy and no brain metastases. In real-world practice, melanoma combination ICI is often used in ECOG 2 patients, those with active brain metastases (increasingly eligible under

expanded label), and those who received prior adjuvant ICI. Each of these groups has substantially worse outcomes than the trial population. The implication is not that combination ICI does not work in these groups -- the HR 0.62 in real-world still represents meaningful benefit -- but that clinicians extrapolating the trial HR 0.45 to clinical decisions overestimate the absolute benefit for individual patients. Real-world evidence corrects this optimism bias with population-representative effectiveness data.

5.2 HTE Changes Practice for Bevacizumab and CDK4/6i

The causal forest HTE findings have direct prescribing implications that average treatment effect analysis would have missed. The finding that CRC bevacizumab benefit was limited to RAS/RAF wild-type patients (HR 0.72 vs 0.96 in mutant) provides population-level evidence for what has been clinically suspected but incompletely characterised in trials. For CDK4/6i in breast cancer, the de novo metastatic benefit (HR 0.48) versus relapsed benefit (HR 0.62) provides a rationale for prioritising CDK4/6i in de novo patients if access or cost constraints require treatment prioritisation. These insights illustrate the clinical value of moving beyond 'does this treatment work on average' to 'which patients benefit most' -- a question that HTE analysis answers and standard outcomes research cannot.

5.3 Limitations

Three limitations. First, despite ML-enhanced confounding control, unmeasured confounders (physician preference, patient values, informal performance status assessment) cannot be fully addressed by any observational method. Second, molecular marker completeness at 74.8% introduces selection bias if missingness is related to treatment or outcome. Third, the HTE findings from causal forests are exploratory and require prospective validation in stratified trials before informing treatment guidelines.

6. Conclusion

6.1 Key Findings

ML-enhanced causal inference achieves the highest Oncology Outcomes Research Score (0.930), closely replicating trial efficacy estimates (2.8-3.6% deviation for NSCLC and breast cancer) while revealing clinically important HTE. The melanoma trial-to-RWE gap (38%) reflects the exclusion of ECOG 2+ and brain metastasis patients from pivotal trials. Bevacizumab benefit in

CRC is limited to RAS wild-type patients in real-world populations. Clinical registries (OORS 0.916) provide the most complete molecular and staging data for confounding adjustment and the highest trial-RWE replication consistency.

6.2 Recommendations

Three recommendations. First, oncology treatment decisions should be informed by real-world effectiveness evidence from registry-based studies alongside trial efficacy data, with the real-world evidence providing population-representative effectiveness estimates for the patient groups most frequently excluded from trials (ECOG 2+, elderly, multimorbid). Second, regulatory agencies should require HTE analysis using causal forest or equivalent methods in oncology real-world evidence submissions, moving beyond average treatment effects to identify the patients most likely to benefit. Third, European cancer registries should harmonise their molecular marker and performance status data collection to enable comparative real-world evidence across national borders, building a pan-European oncology RWD platform with the statistical power to address rare subgroup questions. All data, causal models, and the OORS calculator are available from the corresponding author on request.

References

- Abadie, A. et al. (2010). Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program. *Journal of the American Statistical Association*, 105(490), 493-505.
- Booth, C. M. et al. (2019). Real-world data about patients with cancer who receive systemic therapy. *Journal of Clinical Oncology*, 37(4), 3060-3067.
- Griffith, S. D. et al. (2019). Using electronic health record-based real-world evidence to complement clinical trial design and analysis. *npj Digital Medicine*, 2(1), 116.
- Khozin, S. et al. (2017). Real-world characterization of patients with advanced non-small cell lung cancer and nivolumab therapy using a novel insurance claims-linked database. *Journal of Clinical Oncology*, 35(Suppl), e20623.
- Miksad, R. A. et al. (2018). Harnessing the power of real-world evidence for small populations. *Nature Reviews Drug Discovery*, 18(9), 688-698.
- Rosenbaum, P. R. and Rubin, D. B. (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1), 41-55.
- Sherman, R. E. et al. (2016). Real-world evidence: What is it and what can it tell us? *New England Journal of Medicine*, 375(23), 2293-2297.

- Unger, J. M. et al. (2021). Patient eligibility and enrollment in clinical trials: Recent findings. ASCO Educational Book, 40, e25-e35.
- Vansteelandt, S. and Keiding, N. (2011). Invited commentary: G-computation: Lost in translation? American Journal of Epidemiology, 173(7), 739-742.
- Wager, S. and Athey, S. (2018). Estimation and inference of heterogeneous treatment effects using random forests. Journal of the American Statistical Association, 113(523), 1228-1242.
- Blumenthal, G. M. et al. (2020). Oncology drug approvals: Evaluating endpoints and evidence in an era of breakthrough therapies. Oncologist, 22(7), 762-767.
- Corrigan-Curay, J. et al. (2018). Real-world evidence and real-world data for evaluating drug safety and effectiveness. JAMA, 320(9), 867-868.
- Flaherty, K. T. et al. (2020). Molecular landscape and actionable alterations in a genomically guided cancer clinical trial. Journal of Clinical Oncology, 38(33), 3883-3894.
- Hernan, M. A. and Robins, J. M. (2016). Using big data to emulate a target trial when a randomized trial is not available. American Journal of Epidemiology, 183(8), 758-764.
- Kent, D. M. et al. (2020). The Predictive Approaches to Treatment effect Heterogeneity (PATH) statement. Annals of Internal Medicine, 172(1), 35-45.
- Lash, T. L. et al. (2021). Good practices for quantitative bias analysis. International Journal of Epidemiology, 43(6), 1969-1985.
- Leyrat, C. et al. (2019). Propensity score analysis with partially observed covariates: How should multiple imputation be used? Statistical Methods in Medical Research, 28(1), 3-19.
- Maringe, C. et al. (2020). The impact of the COVID-19 pandemic on cancer deaths due to delays in diagnosis in England, UK. The Lancet Oncology, 21(8), 1023-1034.
- Mok, T. S. K. et al. (2018). Pembrolizumab versus chemotherapy for previously untreated, PD-L1-expressing, locally advanced or metastatic non-small-cell lung cancer (KEYNOTE-042). The Lancet, 393(10183), 1819-1830.
- Salazar, M. C. et al. (2019). Association of intensity of care at end of life with survival for patients with lung cancer. JAMA Oncology, 5(9), 1240-1246.

Declarations

Funding

Supported by the French National Research Agency (ANR-23-CE36-0004) and the Italian Ministry of Health (Ricerca Corrente 2023). No funder influenced study design, data analysis, or publication.

Conflict of Interest

Both authors declare no competing interests. Neither has received funding from any pharmaceutical company or real-world data platform provider.

Data Availability Statement

Cancer registry data access is governed by national data access committees. Aggregated study results and the OORS calculator are available from the corresponding author. ML causal inference code (TMLE + causal forest pipeline) is available on request.

Ethical Approval

Approved by Institut Gustave Roussy Ethics Committee (IGR-2023-248), INT Milan Comitato Etico (No. 228-2023), Karolinska Regional Ethics (Dnr 2023-04486), and Charite Ethics Committee (EA4/462/23). Access to patient-level registry data was approved by national data access bodies.

Appendix A

OORS Methodology and ML Causal Inference Pipeline

This appendix details the OORS scoring rubric, the TMLE implementation, and the causal forest HTE estimation protocol.

Part I -- OORS Dimensions

- 1a. Population Representativeness: Demographic + clinical similarity to cancer registry reference Standardised difference on 12 key variables
- 1b. Causal Validity: E-value analysis + covariate balance (max SMD) + sensitivity analyses
Higher E-value + lower SMD = stronger causal claim
- 2a. Data Completeness: Fraction of key oncology confounders available Stage, molecular markers, PS, prior treatment
- 2b. Policy Actionability: Payer, regulator, or guideline uptake within 24 months Expert panel rating
- 3a. Replication Consistency: Deviation of RWE estimates from available trial estimates
Weighted by trial quality
- 3b. $OORS = 0.20 * Represent. + 0.25 * Causal + 0.20 * Data + 0.20 * Policy + 0.15 * Replication$

Part II -- ML Causal Inference Pipeline

- 4a. Propensity model: SuperLearner ensemble (lasso, gradient boosting, random forest)
Cross-validated, stabilised IPTW
- 4b. Outcome model: SuperLearner with same learners G-computation for potential outcomes
- 5a. TMLE: One-step targeted correction for efficient doubly-robust estimation 95% CI via influence function
- 5b. HTE: Causal forest (R grf package), 2000 trees, honest splitting Variable importance + partial dependence plots