Paper review:
A nonparametric framework for treatment effect modifier discovery in high dimensions

Philippe Boileau, Ning Leng, Nima S. Hejazi, Mark van der Laan and Sandrine Dudoit (2025)

Suehyun Kim

3 February 2026

Causal Inference Lab.
Seoul National University

Table of contents

Table of contents

Introduction

In causal inference, discovering **treatment effect modifiers (TEMs)** is crucial for understanding heterogeneous treatment effects in a population.

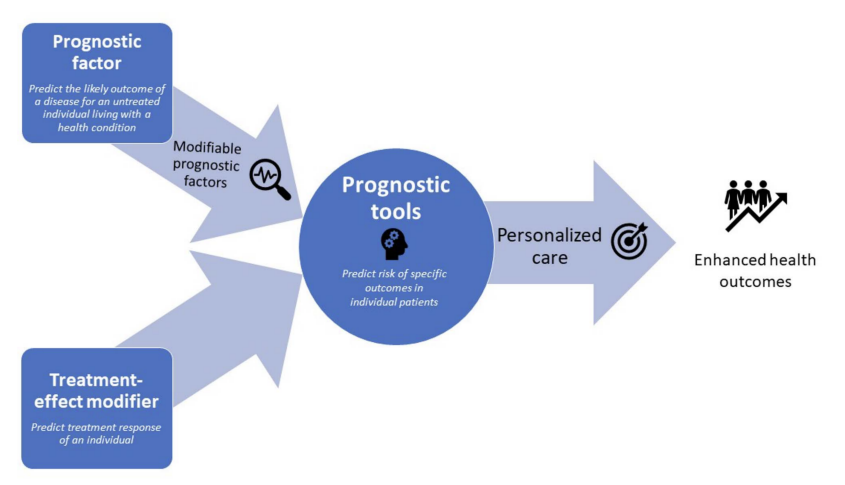

Figure reference: Tousignant-Laflamme et al. (2022)

- Precision medicine: Identifying patient subgroups exhibiting different benefits from a therapy
- Economics: Assessing government policies on diverse population strata

Prior approaches to TEM discovery.

- **Parametric modeling**: Add treatment-covariate interaction terms in a linear outcome model, and assess the statistical significance of interaction coefficients.
- **CATE estimation**: Assess variable imporatance in CATE prediction.
  - Permutation-based / Dropout-based approaches
  - Penalized regression methods with feature selection properties (e.g. LASSO)

However, the above methods *rely on unverifiable assumptions* (model misspecification, sparsity, correlation structures) or produce *unreliable results in high-dimensional settings*.

Introduction

In this paper, we explore a general framework for defining and performing inference about marginal **TEM variable importance parameters (TEM-VIPs)**.

- This work generalizes the method of Boileau et al. (2022) to continuous and binary outcomes, which directly estimates the strength of covariates' capacity to modify the treatment effect.
- The proposed framework does not rely on stringent assumptions on the DGP and accommodates high-dimensional settings.
- The framework is equipped with tools to define appropriate TEM-VIPs and corresponding EIF-based nonparametric estimators that allow valid and reliable inference.

Table of contents

Background on nonparametric asymptotic theory

Notation.

- **Full data**: $\{X_i\}_{i=1}^n$, $X_i = (W_i, A_i, Y_i^{(0)}, Y_i^{(1)}) \overset{\text{i.i.d.}}{\sim} P_{X,0} \in \mathcal{M}_X$
  - W_i: Set of p covariates, possibly $p \gg n$
  - A_i: Binary treatment indicator
  - $Y_i^{(a)}$: Potential outcomes; assume $Y_i^{(a)} \in (0, 1)$
  - $\mathcal{M}_X$: Nonparametric model of possible DGPs
- **Observed data**: $O = (W, A, Y) \sim P_0 \in \mathcal{M}$
  - Consistency: $Y = AY^{(1)} + (1 - A)Y^{(0)}$
  - Under consistency, $\mathcal{M}$ is fully determined by $\mathcal{M}_X$

Background on nonparametric asymptotic theory

Recall that $P(Y, A, W) = P(Y \mid A, W)\, P(A \mid W)\, P(W)$. Thus, we define the following nuisance functions:

- P_0: DGP of observed data
- P_n: Empirical distribution
- $\hat{P}_n$: Plug-in estimator of P_n, made up of elements of P_n and possibly nuisance parameter estimators
- Expected conditional outcome: $\bar{Q}_0(A, W) = \mathbb{E}_{P_0}[Y \mid A, W]$
  - Full-data counterpart: $\bar{Q}_{P_{X,0}}(a, W) = \mathbb{E}_{P_{X,0}}[Y^{(a)} \mid W]$
- Propensity score: $g_0(W) = P_{P_0}[A = 1 \mid W]$

Background on nonparametric asymptotic theory

Nonparametric EIF-based estimators possess numerous attractive properties:

- Efficiency
- Root-n consistency and asymptotic normality
- No assumptions on outcome/PS models (robustness, flexibility)

Background on nonparametric asymptotic theory

Motivation. Let $\Theta(P)$ be the parameter of interest and $\Theta(\hat{P}_n)$ be its plug-in estimator.

- Let $D(O, P)$ be the influence function $\Theta(P)$ at P.
- The asymptotic behavior of $\Theta(\hat{P}_n)$ is captured by the following von Mises expansion:

$$\sqrt{n}\big(\Theta(\hat{P}_n) - \Theta(P_0)\big) = \underbrace{\sqrt{n}\, E_{P_n}[D(O, P_0)]}_{\text{Term 1}} - \underbrace{\sqrt{n}\, E_{P_n}\Big[D(O, \hat{P}_n)\Big]}_{\text{Term 2}} + \underbrace{\sqrt{n}\,(E_{P_n} - E_{P_0})\Big(D(O, \hat{P}_n) - D(O, P_0)\Big)}_{\text{Term 3}} - \underbrace{\sqrt{n}\, R(P_0, \hat{P}_n)}_{\text{Term 4}}.$$

Background on nonparametric estimation theory

- Term 1: $\sqrt{n}\, E_{P_n}[D(O, P_0)]$
  - Converges to a mean-zero Gaussian r.v. with variance $\mathbb{E}_{P_0}[D(O, P_0)^2]$
  - **EIF-based estimators minimize this asymptotic variance**
- Term 2: $\sqrt{n}\, E_{P_n}\left[D(O, \hat{P}_n)\right]$
  - Asymptotic bias
  - We construct estimators that eliminate this bias term
- Terms 3 & 4: Converges to 0 under standard assumptions

Background on nonparametric estimation theory

Estimators. By leveraging the EIF $D(O, P)$, nonparametric estimators correct first-order bias and attain semiparametric efficiency.

(i) **One-step**: Adds empirical EIF correction

$$\Theta^{(\mathrm{OS})}(\hat{P}_n) = \Theta(\hat{P}_n) + \mathbb{E}_{P_n}[D(O, \hat{P}_n)]$$

(ii) **Estimating equation**: Solves the estimating equation

$$\mathbb{E}_{P_n}[D(O, \hat{P}_n)] = 0.$$

(iii) **Targeted maximum likelihood (TML)**: Tilt $\hat{P}_n$ to generate $P_n^\star$ such that $\mathbb{E}_{P_n}[D(O, \hat{P}_n^\star)] \approx 0$, and define $\Theta^{\mathrm{TML}}(\hat{P}_n) = \Theta(P_n^\star)$.

- TML estimators constrain estimates to the parameter space.
- We will look at an example later.

Inference. Inference about $\Theta(P_0)$ follows naturally from *asymptotic normality*.

In particular, the α-level Wald-type confidence interval for $\Theta(P_0)$ can be constructed identically for the three estimators $\Theta^{(*)}(\hat{P}_n)$:

$$\Theta^{(*)}(\hat{P}_n) \pm z_{1-\alpha/2}\sqrt{\frac{\mathbb{E}_{P_0}[D(O,P_0)^2]}{n}},$$

where the term $\mathbb{E}_{P_0}[D(O,P_0)^2]$ is substituted with $\mathbb{E}_{P_0}[D(O,\hat{P}_n)^2]$ in practice.

Example: EIF of ATE

Intuitively, the EIF is the *orthogonal projection* of any influence function onto the *tangent space of the observed-data model*.

- Influence function is the analogue of a *gradient vector* (in standard calculus).
- Consider, for instance, the ATE defined in terms of the full-data model

$$\Theta^F(P_{X,0}) = \mathbb{E}_{P_{X,0}}[Y^{(1)} - Y^{(0)}].$$

- If we had access to full data, the influence function would be given by $Y^{(1)} - Y^{(0)} - \Theta(P_{X,0})$.
- We project this IF onto the tangent space of the observed-data model to obtain

$$\begin{aligned} D(O, P_0) = &\frac{A}{g_0(W)}(Y - \bar{Q}_0(1, W)) - \frac{1 - A}{1 - g_0(W)}(Y - \bar{Q}_0(0, W)) \\ &+ \bar{Q}_0(1, W) - \bar{Q}_0(0, W) - \Theta(P_0). \end{aligned}$$

Example: EIF of ATE

Deriving one-step and estimating equation estimators of the ATE using the EIF, we obtain

$$\Theta_{\text{OS}} = \Theta_{\text{EE}} = \frac{1}{n}\sum_{i=1}^{n}\left[\left(\frac{A_i}{g_0(W_i)} - \frac{1-A_i}{1-g_0(W_i)}\right)\{Y_i - \bar{Q}_0(A_i, W_i)\} + \bar{Q}_0(1, W_i) - \bar{Q}_0(0, W_i)\right].$$

which coincides with the AIPW estimator for the ATE.

Table of contents

TEM-VIPs for continuous outcomes

Causal parameter. Suppose the outcome is continuous, and the treatment effect of interest is the ATE, i.e. $\mathbb{E}_{P_{X,0}}[Y^{(1)} - Y^{(0)}]$.

First, we impose assumptions for computational convenience:

(A1) Centered covariates: $\mathbb{E}_{P_{X,0}}[W_j] = 0$.

(A2) Nonzero variance: $\mathbb{E}_{P_{X,0}}[W_j^2] > 0$.

TEM-VIPs for continuous outcomes

Definition (TEM-VIP, with respect to ATE)

An *(absolute) TEM-VIP* of the j-th covariate is defined as a mapping

$$\begin{aligned}\Psi_j^F(P_{X,0}) &:= \frac{\mathrm{Cov}_{P_{X,0}}\left[Y^{(1)} - Y^{(0)},\, W_j\right]}{\mathbb{V}_{P_{X,0}}\left[W_j\right]} \\ &= \frac{\mathbb{E}_{P_{X,0}}\left[(Y^{(1)} - Y^{(0)})W_j\right]}{\mathbb{E}_{P_{X,0}}\left[W_j^2\right]} \\ &= \frac{\mathbb{E}_{P_{X,0}}\left[\left(\bar{Q}_{P_{X,0}}(1,W) - \bar{Q}_{P_{X,0}}(0,W)\right) W_j\right]}{\mathbb{E}_{P_{X,0}}\left[W_j^2\right]}.\end{aligned}$$

The full estimand is then given by

$$\Psi^F : \mathcal{M}_X \to \mathbb{R}^p, \qquad \Psi^F(P_{X,0}) = \left(\Psi_1^F(P_{X,0}), \ldots, \Psi_p^F(P_{X,0})\right).$$

TEM-VIPs for continuous outcomes

Interpretation. The TEM-VIP $\Psi^F(P_{X,0})$ can be viewed as assessing the *correlation between the difference in potential outcomes and each potential TEM*, renormalized to be on the same scale as Y.

- Assume that the expectation of $f(W) = \bar{Q}_{P_{X,0}}(1, W) - \bar{Q}_{P_{X,0}}(0, W)$ conditioned on W_j is linear in W_j, i.e. $\mathbb{E}_{P_{X,0}}[f(W) \mid W_j] = \beta_j W_j$. Then, $\Psi^F(P_{X,0})$ is the vector of *simple linear regression coefficients*:

$$\Psi_j^F(P_{X,0}) = \frac{\mathbb{E}_{P_{X,0}}[f(W)W_j]}{\mathbb{E}_{P_{X,0}}[W_j^2]} = \frac{\mathbb{E}_{P_{X,0}}[\beta_j W_j^2]}{\mathbb{E}_{P_{X,0}}[W_j^2]} = \beta_j.$$

- When the relationship between $f(W)$ and the W_j's are nonlinear, the parameter can be interpreted as a linear model projection:

$$\Psi_j^F(P_{X,0}) = \underset{\beta_j \in \mathbb{R}}{\operatorname{argmin}}\, \mathbb{E}_{P_{X,0}}\left[\left(\left(\bar{Q}_{P_{X,0}}(1, W) - \bar{Q}_{P_{X,0}}(0, W)\right) - (\alpha + \beta_j W_j)\right)^2\right].$$

TEM-VIPs for continuous outcomes

Under *(A3) unconfoundedness* and *(A4) positivity* assumptions, the full-data TEM-VIP $\Psi_j^F(P_{X,0})$ is identifiable from observed data:

Theorem 1 (Identifiability of TEM-VIP)

Assuming that (A1), (A2), (A3) and (A4) hold, we have

$$\Psi_j(P_0) := \frac{\mathbb{E}_{P_0}\left[\left(\bar{Q}_0(1,W) - \bar{Q}_0(0,W)\right) W_j\right]}{\mathbb{E}_{P_0}\left[W_j^2\right]} = \Psi_j^F(P_{X,0}).$$

Hence, the parameter $\Psi : \mathcal{M} \to \mathbb{R}^p$, $\Psi(P_0) = (\Psi_1(P_0), \ldots, \Psi_p(P_0))$ is equal to the full-data estimand $\Psi^F(P_{X,0})$.

TEM-VIPs for continuous outcomes

Efficient influence function. The EIF provides the basis for the construction of nonparametric estimators.

Proposition 1 (EIF of observed-data parameter)

Under (A1) and (A2), the EIF at $P \in \mathcal{M}$ of $\Psi_j(P)$ is given by

$$D_j(O,P) = \frac{W_j}{\mathbb{E}_P\left[W_j^2\right]}\left(\frac{2A-1}{Ag(W)+(1-A)(1-g(W))}\left(Y-\bar{Q}(A,W)\right)\right.$$
$$\left.+\bar{Q}(1,W)-\bar{Q}(0,W)-\Psi_j(P)\,W_j\right).$$

Remark. Note that the blue term corresponds to the *uncentered EIF* of the ATE, which is the treatment effect of interest.

Estimators.

(i) One-step and estimating equation estimators: The one-step and estimating equation estimators of $\Psi_j(P_0)$ are identical. Let $\bar{Q}_n$ and g_n be estimators of $\bar{Q}_0$ and g_0 trained of P_n and included in $\hat{P}_n$. Then,

$$\Psi_j^{(\mathrm{OS})}(\hat{P}_n) = \Psi_j^{(\mathrm{EE})}(\hat{P}_n) = \frac{1}{\sum_{i=1}^n W_{ij}^2} \sum_{i=1}^n W_{ij} \times \left(\frac{2A_i - 1}{A_i g_n(W_i) + (1 - A_i)(1 - g_n(W_i))} \left(Y_i - \bar{Q}_n(A_i, W_i)\right) + \bar{Q}_n(1, W_i) - \bar{Q}_n(0, W_i) \right).$$

(ii) Targeted maximum likelihood (TML) estimator

Step 1. Define the negative log-likelihood loss function for $\bar{Q}$ as

$$L(O; \bar{Q}) = -\log\left\{ \bar{Q}(A, W)^Y \left(1 - \bar{Q}(A, W)\right)^{1-Y} \right\},$$

and a parametric working submodel for $\bar{Q}$ as

$$\bar{Q}_j(\epsilon)(A, W) = \text{logit}^{-1}\left\{ \text{logit}\, \bar{Q}(A, W) + \epsilon H_j(A, W) \right\},$$

where

$$H_j(A, W) := \frac{W_j}{E_P[W_j^2]} \frac{2A - 1}{Ag(W) + (1 - A)(1 - g(W))}.$$

TEM-VIPs for continuous outcomes

Step 2. Denoting an initial estimator of $\bar{Q}_0$ trained on P_n by $\bar{Q}_n^0$, update $\bar{Q}_n^0$ by computing $\epsilon_{n,j}^1$ such that

$$\epsilon_{n,j}^1 = \arg\min_{\epsilon} E_{P_n}\left[L\left(O; \bar{Q}_{n,j}^0(\epsilon)\right)\right].$$

Compute the *tilted conditional outcome estimator* $\bar{Q}_{n,j}^1 = \bar{Q}_{n,j}^0(\epsilon_{n,j}^1)$ and the *tilted distribution* $P_n^\star$, where $\bar{Q}_n^0$ is replaced by $\bar{Q}_{n,j}^1$ in $\hat{P}_n$.

Step 3. The TML estimator of the j-th TEM-VIP is given by

$$\Psi_j^{\mathsf{TML}}(\hat{P}_n) = \Psi_j(P_n^\star).$$

TEM-VIPs for continuous outcomes

Asymptotic behavior. Note that the asymptotic distributions of $\Phi^{\text{OS}}(\hat{P}_n)$, $\Phi^{\text{EE}}(\hat{P}_n)$ and $\Phi^{\text{TML}}(\hat{P}_n)$ are identical due to their dependence on the common EIF $D_j(O, P_0)$.

First, we establish *consistency* results and its double-robustness.

(A5) Outcome model estimator consistency: $\|\bar{Q}_n(A, W) - \bar{Q}_0(A, W)\|_2^2 = o_P(1)$.

(A6) PS estimator consistency: $\|g_n(W) - g_0(W)\|_2^2 = o_P(1)$.

Proposition 2 (Consistency)

Under (A1) and (A2), and *either (A5) or (A6)*,

$$\Phi^{(*)}(\hat{P}_n) \xrightarrow{P} \Phi(P_0), \qquad * = \text{OS, EE, TML.}$$

TEM-VIPs for continuous outcomes

Moreover, the *asymptotic normality* of the estimators' smapling distributions can be specified under the following assumptions:

(A7) Donsker conditions: There exists a P_0-Donsker class $\mathcal{G}_0$ such that
$P_{P_0}\left[D_j(O, \hat{P}_n) \in \mathcal{G}_0\right] \to 1$ and
$E_{P_0}\left[\left(D_j(O, \hat{P}_n) - D_j(O, P_0)\right)^2 \mid \hat{P}_n\right] = o_P(1)$ for each j.

(A8) Shared rate convergence:
$\|\bar{Q}_n(A, W) - \bar{Q}_0(A, W)\|_2 \, \|g_n(W) - g_0(W)\|_2 = o_P(n^{-1/2})$.

(A9) Uniformly bounded covariates: $|W_j| \leq C$ for $j = 1, \ldots, p$.

Theorem 2 (Asymptotic normality)

Under (A1), (A2), (A7), (A8) and (A9),

$$\sqrt{n}\left(\Psi_j^{(*)}(\hat{P}_n) - \Psi_j(P_0)\right) \xrightarrow{D} N\left(0, E_{P_0}[D_j(O, P_0)^2]\right), \quad * = \text{OS, EE, TML.}$$

Table of contents

General framework for TEM-VIPs

The workflow presented in the previous section for continuous outcomes can be unified into a general framework.

For instance, suppose we are interested in causal effects on *right-censored time-to-event* outcomes, which are common in clinical trials.

- Causal parameters of interest are often built upon the conditional survival function $S_{P_{X,0}}(t \mid a, W) := P_{P_{X,0}}[T^{(a)} > t \mid W]$, where $T^{(a)}$ denotes the (potential) event time.
- **CATE of survival probability at time t**:

$$\mathbb{E}_{P_{X,0}}\left[S_{P_{X,0}}(t \mid 1, W) - S_{P_{X,0}}(t \mid 0, W) \,\middle|\, W\right].$$

- **Difference in conditional restricted mean survival times (RMSTs)**:

$$\mathbb{E}_{P_{X,0}}\Big[\min\{T^{(1)}, t\} - \min\{T^{(0)}, t\} \; W = \mathbb{E}_{P_{X,0}}\left[\int_0^t \left\{S_{P_{X,0}}(u \mid 1, W) - S_{P_{X,0}}(u \mid 0, W)\right\} du \,\middle|\, W\right].$$

General framework for TEM-VIPs

Step 1. Select a full-data, pathwise differentiable parameter $\Phi^F(P_{X,0})$ of some treatment effect that is relevant to the problem at hand.

- ATE (for continuous/binary outcomes)

$$\Phi^F(P_{X,0}) = \mathbb{E}_{P_{X,0}}[Y^{(1)} - Y^{(0)}]$$

- Difference in RMSTs (for right-censored time-to-event outcomes)

$$\Phi^F(P_{X,0}) = \mathbb{E}_{P_{X,0}}[\min\{T^{(1)}, t\} - \min\{T^{(0)}, t\}]$$

General framework for TEM-VIPs

Step 2. Define $f(W)$ such that $\mathbb{E}_{P_{X,0}}[f(W)] = \Phi^F(P_{X,0})$.

- CATE (for estimation of ATE)

$$f(W) = \bar{Q}_{P_{X,0}}(1, W) - \bar{Q}_{P_{X,0}}(0, W)$$

- Conditional difference in RMSTs (for difference in RMSTs)

$$f(W) = \int_0^t \left\{ S_{P_{X,0}}(u \mid 1, W) - S_{P_{X,0}}(u \mid 0, W) \right\} \mathrm{d}u$$

Under (A1) and (A2), the **TEM-VIP** of the j-th covariate Θ_j^F is given by

$$\Theta_j^F = \frac{\mathbb{E}_{P_{X,0}}[f(W) W_j]}{\mathbb{E}_{P_{X,0}}[W_j^2]}.$$

General framework for TEM-VIPs

Step 3. Establish the *identifiability* of the TEM-VIP.

- Let Θ_j and Φ be the observed-data counterparts of Θ_j^F and Φ^F, respectively.
- The conditions establishing $\Theta_j^F(P_{X,0}) = \Theta_j(P_0)$ are identical to the conditions for the identifiability of the treatment effect of interest, i.e., $\Phi^F(P_{X,0}) = \Phi(P_0)$.
  - e.g., consistency, unconfoundedness, positivity.
  - The only additional assumption is that W_j has bounded variance.

General framework for TEM-VIPs

Step 4. Derive the *EIF* of the TEM-VIP.

If the uncentered EIF of $\Phi(P_0)$ is given by $d(O, P_0)$, then the EIF of the TEM-VIP is given by

$$\frac{W_j}{\mathbb{E}_{P_0}[W_j^2]} \left\{ d(O, P_0) - W_j \Theta_j(P_0) \right\}.$$

General framework for TEM-VIPs

Step 5. Construct the EIF-based estimators for the TEM-VIP.

- These include *(i) one-step*, *(ii) estimating equation* and *(iii) TML* estimators.
- The asymptotic properties of the EIF-based estimators are identical to those of the nonparametric efficient estimators of Φ, such as double-robustness, provided that W_j is bounded.

Table of contents

Simulation studies

Through simulation studies, we observe:

(i) *Finite-sample performance* of the proposed one-step and TML estimators.

(ii) The *estimators' capacity to recover TEMs*, compared to (augmented) modified covariates methods (Tian et al. (2014) & Chen et al. (2017).

We consider three DGPs: *(i) continuous outcome, observational study*, *(ii) binary outcome, observational study* and *(iii) survival outcome, RCT*.

- Sample sizes: $n = 125, 250, 500, 1{,}000, 2{,}000$
- 200 MC simulations under each setting

Simulation setup

DGP 1. Continuous outcome, observational study

$$W \sim N(0, I_{500\times 500})$$

$$A \mid W \sim \text{Bernoulli}\left(\text{logit}^{-1}\left(\frac{1}{4}(W_1 - W_2 + W_3)\right)\right)$$

$$Y \mid A, W \sim 1 + 2\left|\sum_{j=1}^{5} W_j\right| + (5A - 2)\sum_{j=1}^{5} W_j + \varepsilon, \quad \varepsilon \sim N(0, 1/2)$$

DGP 2. Binary outcome, observational study

$$W \sim N(0, \Sigma_{100\times 100}), \quad \Sigma_{ij} = \begin{cases} 1, & i = j, \\ 0.1|i-j|^{-1.8}, & \text{otherwise} \end{cases}$$

$$A \mid W \sim \text{Bernoulli}\left(\text{logit}^{-1}\left(\frac{1}{4}(W_1 + W_2 + W_3)\right)\right)$$

$$Y \mid A, W \sim \text{Bernoulli}\left(\text{logit}^{-1}\left(1 - 2A + \sum_{j=1}^{5} W_j + \left(A - \frac{1}{2}\right)\sum_{j=1}^{5} W_j\right)\right)$$

Simulation setup

DGP 3. Right-censored time-to-event outcome, randomized control trial

$$
\begin{aligned}
W &\sim N(0, \Sigma_{300\times 300}) \\
A &\sim \text{Bernoulli}\left(\tfrac{1}{2}\right) \\
C \mid A, W &\sim \min\left\{\text{NegBin}\left(1,\, \text{logit}^{-1}(5 + A + W_1)\right),\, 10\right\} \\
T \mid A, W, C &\sim \text{NegBin}\left(1,\, \text{logit}^{-1}\left(-2 - A + (10A - 5)\sum_{j=1}^{10} W_j\right)\right) \\
\tilde{T} &= \min\{T, C\} \\
\Delta &= \mathbb{I}(T > c)
\end{aligned}
$$

- The covariance matrix Σ is block-diagonal, with each block corresponding to ten moderately correlated features.
- The target estimand is defined based on difference in RMSTs at $t = 9$ (total duration consists of 10 time units).

Simulation 1: Finite-sample performance

Simulation 1. Finite-sample performance of estimators

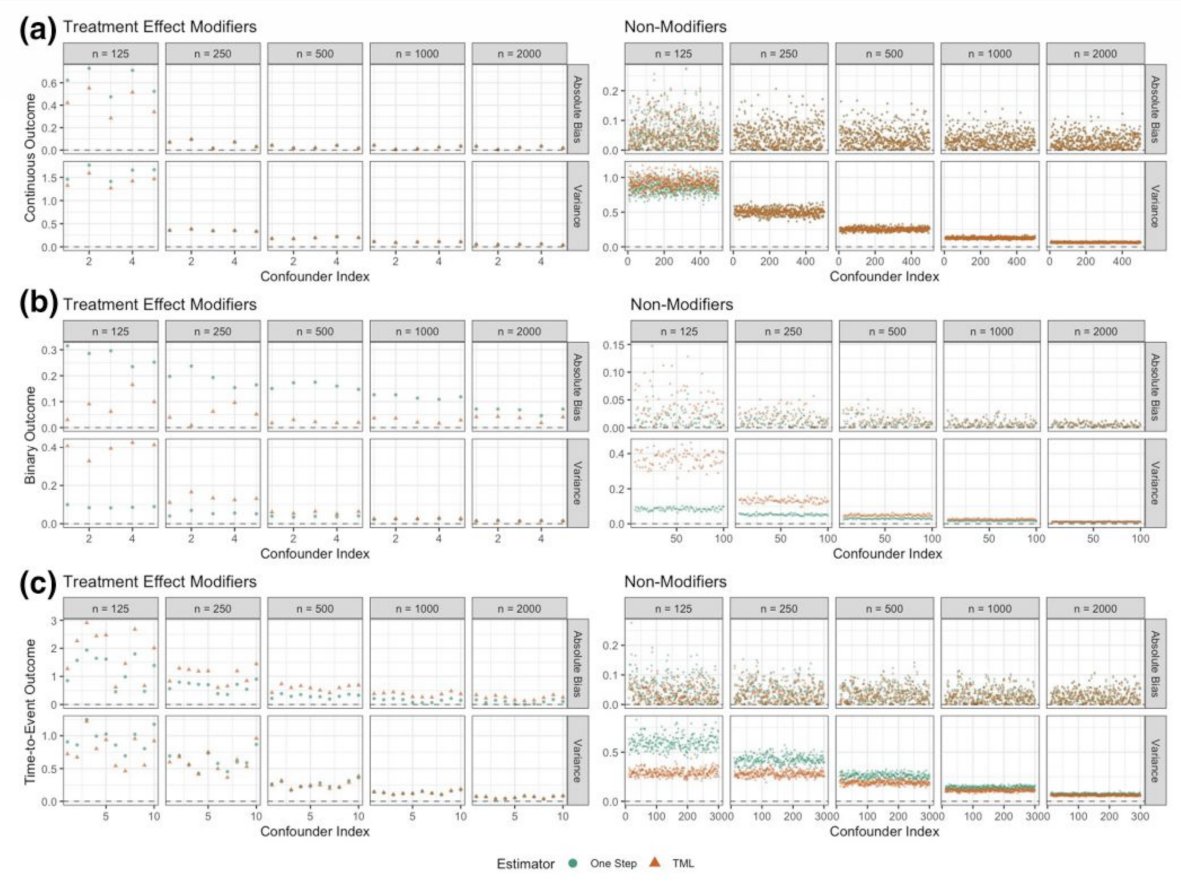

Figure 1. Empirical bias and variance of one-step and TML estimators in the continuous, binary, and time-to-event outcome simulation scenarios.

Simulation 2: TEM discovery and classification

Simulation 2. Estimators' capacity to recover TEMs

- The estimators' ability to distinguish TEMs is evaluated in terms of empirical false discovery rate (FDR), true negative rate (TNR) and true positive rate (TPR).
- Two-sided Wald-type tests at the 5% level were adjusted for multiple testing using the Benjamini–Hochberg procedure.
- The one-step and TML estimators are compared to modified covariates and augmented modified covariates methods with LASSO.
  - Modified covariates: Outcomes are transformed so that only the treatment-covariate interactions in a GLM need be modeled.
  - Augmented modified covariates: The transformed outcomes are modeled as a function of all covariates to improve efficiency.

Simulation 2: TEM discovery and classification

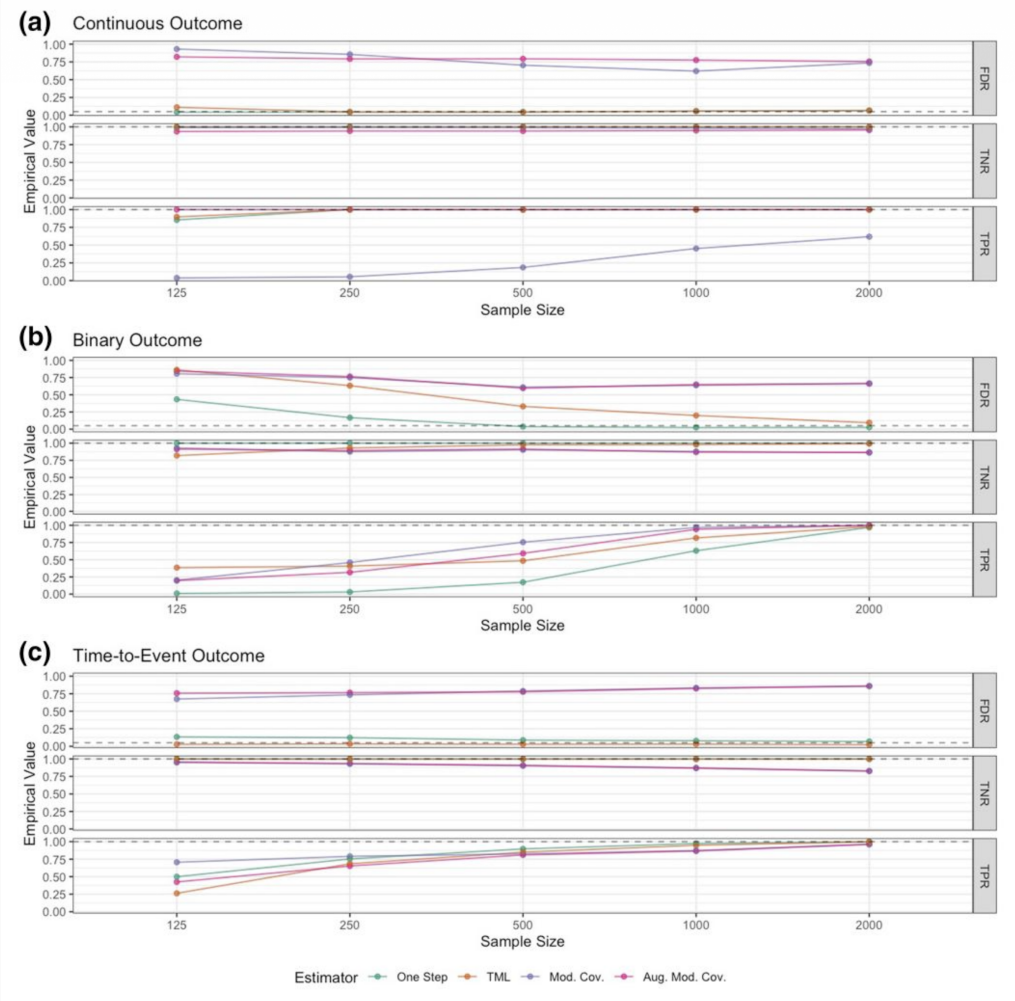

Figure 2. TEM classification results of one-step, TML, modified covariates, augmented modified covariates estimators in the continuous, binary, and time-to-event outcome simulation scenarios.

Table of contents

Data application

The proposed framework is applied to a clinical trial dataset with a *right-censored time-to-event outcome*.

Data source. Subset of FinHER dataset provided by Loi et al. (2014)

- Trastuzumab significantly improves clincal outcomes of breast cancer patients with tumors overexpressing HER2; however, the improvement is not uniform.
- Patients with overexpressed HER2 were randomized to receive either 9 weekly trastuzumab infusions or no trastuzumab as adjuvant treatment for early-stage breast cancer.

Data application

Variables of interest.

- **Treatment**: Trastuzumab (antibody targeting the HER2 oncogene)
- **Covariates (potential TEMs)**: 500 genes selected by variability
- **Outcome**: Distant disease-free survival, i.e. time between randomization and first cancer recurrence or death

Implementation of proposed methodology.

- **Target estimand**: RMST-based TEM-VIP
- **Estimation**: TML estimator based on conditional failure and censoring hazards estimates by a Super Learner

Data application

Result 1. Classification of covariates into TEMs and non-TEMs

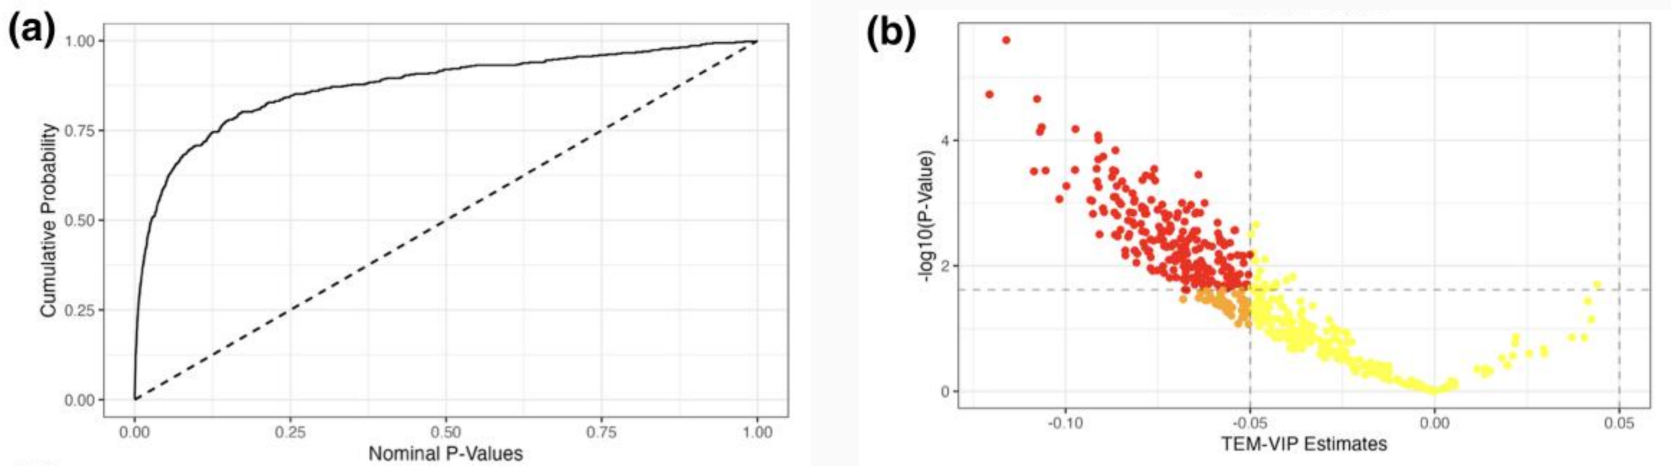

Figure 3. (a) Empirical cdf of p-values.
(b) Volcano plot of the 500 most variable genes' TEM-VIP estimates and associated p-values. Yellow genes are deemed unimportant due to their small estimated effect sizes and larger p-values; orange genes possess a meaningful estimated effect but fail to achieve the adjusted p-value cut-off; red genes are significant at the 5% FDR level and have large estimated TEM-VIPs.

Data application

Result 2. Top five selected treatment effect modifiers

	Gene	Estimate	*SE*	Adj. *p*-value
1	EPPK1	−0.116	0.025	.001
2	NDUFB3	−0.121	0.028	.004
3	BNIP3L	−0.108	0.025	.004
4	PNKD	−0.106	0.027	.006
5	DUSP4	−0.097	0.024	.006

All of the genes above have previously been linked to breast cancer:

- EPPK1: Increased expression has been linked to estrogen-related receptor γ, which is associated with breast cancer growth suppression.
- NDUPF3: A single nucleotide polymorphism in NDUFB3 promoter is reported to be significantly associated with estrogen receptor negative breast cancer.

Data application

Result 3. Clustering of patients with respect to predicted TEMs

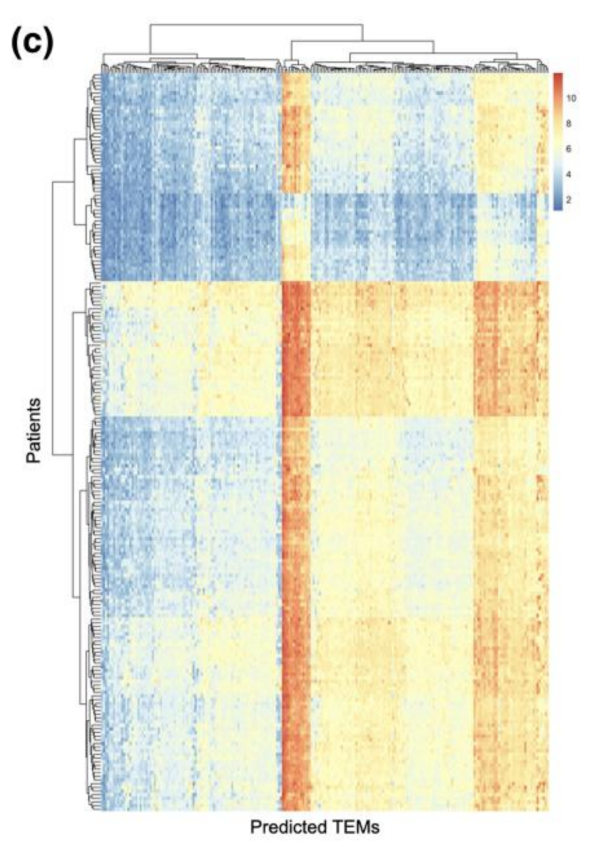

Figure 3. (c) The log-transformed gene expression data of genes with meaningful effect estimates are used to cluster patients. Hierarchical clustering with complete linkage is used for patients and identified TEMs alike.

Data application

Observations.

- *(Result 1)* We observe a continuum in the biomarkers' capacity to influence the treatment effect, in terms of both statistical significance and effect size.
  - Hypothesis testing alone may not be adequate.
  - One can instead deem a biomarker of clnical interest if it is significant at the 5% FDR level and its absolute estimated TEM-VIP is larger than 0.05.
- *(Result 3)* The gene expression data produce multiple distinct patient clusters.
  - However, the authors suggest that it should be considered solely as a diagnostic tool, since redefining subgroups within the same dataset based on the estimated TEM-VIPs may result in overfitting.

Table of contents

Discussion

Contributions of the paper:

(i) Proposed a *general workflow* for defining TEM-VIPs with respect to the treatment effect of interest and for deriving associated *nonparametric estimators based on the EIF*.

(ii) The methodology is applicable to *various types of outcomes*, and the performance is validated under *high-dimensional settings*.

Future directions:

(i) Statistically rigorous subgroup discovery.

(ii) Improving treatment effect estimation in high-dimensional settings by exploiting TEM-VIPs as variable filters.

(iii) Deriving standardized TEM-VIPs.

Reference


Boileau, P., Leng, N., Hejazi, N. S., van der Laan, M. J., & Dudoit, S. (2025). *A nonparametric framework for treatment effect modifier discovery in high dimensions*. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 87(1), 157–185.


Boileau, P., Qi, N. T., van der Laan, M. J., Dudoit, S., & Leng, N. (2022). *A flexible approach for predictive biomarker discovery*. *Biostatistics*, 24(4), 1085–1105.


Gruber, S. and van der Laan, M. J. (2009). *Targeted Maximum Likelihood Estimation: A Gentle Introduction*. *U.C. Berkeley Division of Biostatistics Working Paper Series*. Working Paper 252.