Improving Trial-Based CATE Estimation with Observational Data Borrowing

Samhita Pal, Postdoctoral Researcher, Department of Biostatistics, Vanderbilt University Medical Center

1 Target
Estimate individualized treatment effects in the RCT population

2 Precision
Use large observational datasets to improve efficiency

3 Validity
Retain RCT-based causal identification through calibration rather than direct transportability

4 Mismatch
Handle different covariate sets across sources with MR-OSCAR, CALM, and B-CALM

Main idea: borrow predictive structure from OS, calibrate in the RCT, and improve CATE estimation under shared or mismatched covariates.

RACER: RCT-only CATE estimation via pseudo-outcome regression

How pseudo-outcomes turn individualized treatment-effect estimation into a regression problem

1 Fit arm-specific outcome models in the RCT

Input: randomized trial data

Subjects in RCT (site r)

Covariates	Treatment (randomized)	Outcome
X	$A \in \{-1,1\}$	Y

$A=1$ (Treatment arm)

$A=-1$ (Control arm)

RCT randomization supports causal identification.

2 Construct the CMO (pseudo-outcome mean) from arm-specific fitted means

Arm $A=1$

$\hat{\mu}_1^r(x) = \mathrm{E}[Y \mid X=x, A=1, S=r]$,

Arm $A=-1$

$\hat{\mu}_{-1}^r(x) = \mathrm{E}[Y \mid X=x, A=-1, S=r]$,

Counterfactual mean outcome (CMO)

$$\hat{\mu}^r(x) = \sum_{a \in \{-1,1\}} \pi_{-a}^r(x)\, \hat{\mu}_a^r(x)$$

where $\pi_a^r(x) = \Pr(A=a \mid X=x, S=r)$ are randomization (propensity) probabilities in the RCT.

3 Build pseudo-outcomes and regress to estimate the CATE

Pseudo-outcome for subject i in site r

$$\hat{\psi}_i = \frac{A_i\{Y_i - \hat{\mu}^r(X_i)\}}{\pi_{A_i}^r(X_i)}$$

$$\hat{\tau}_{\mathrm{RACER}} \in \arg\min_{f \in \mathcal{F}} \frac{1}{n^r} \sum_{i: S_i = r} \{\hat{\psi}_i - f(X_i)\}^2$$

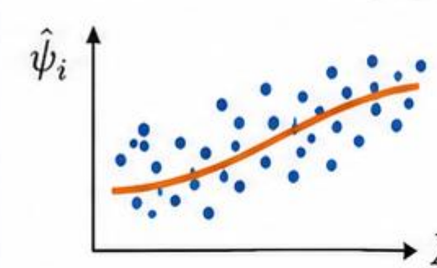

Regress pseudo-outcomes on patient covariates X to learn the CATE surface.

R-OSCAR: penalized calibration of OS outcome models (

Borrow a stable source from the OS, then learn a penalized discrepancy on RCT data only to align with the RCT.

1 Learn from OS (

OS outcome model $\hat{\mu}_a^o(x)$

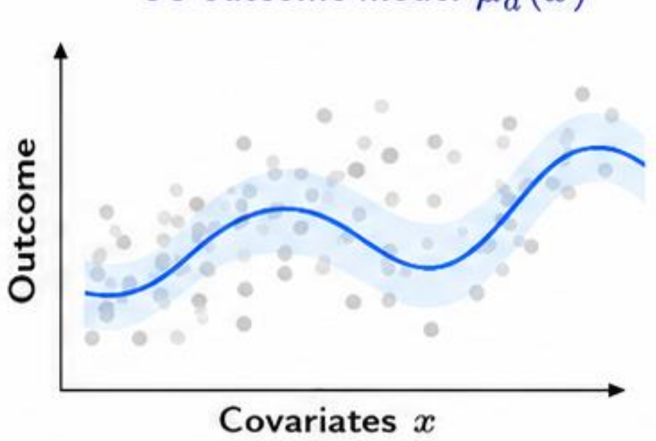

OS prediction $\hat{\mu}_a^o(x)$

2 RCT outcomes vs OS (

RCT data inform discrepancy

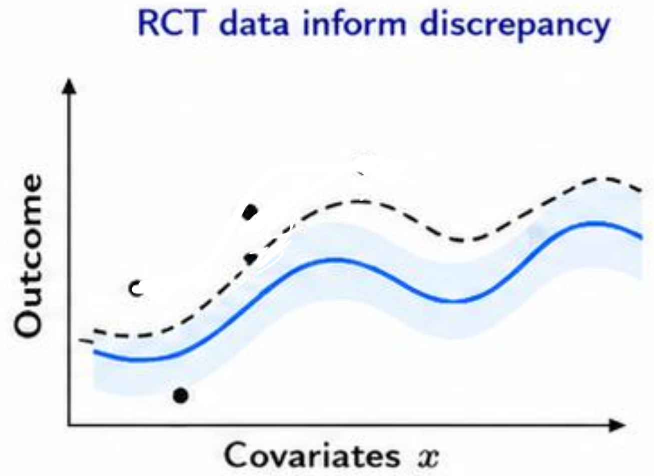

RCT truth (unknown)

3 Learn penalized discrepancy (

Discrepancy $\hat{\delta}_a(x)$ estimated on RCT data only, penalized toward 0

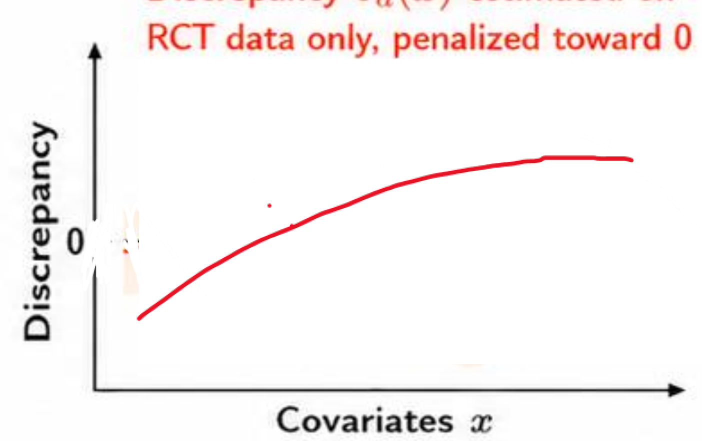

Penalization shrinks toward 0:
- prevents overfitting
- preserves borrowing from OS

4 Calibrated outcome model

$\hat{\mu}_a^{cal}(x) = \hat{\mu}_a^o(x) + \hat{\delta}_a(x)$

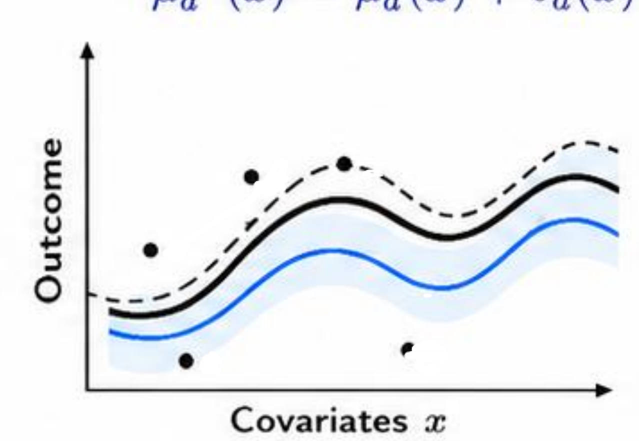

Calibrated model $\hat{\mu}_a^{cal}(x)$

Penalized discrepancy learning for arm a

We learn the discrepancy on the RCT by penalized least squares:

$$\hat{\delta}_a \in \underset{d}{\text{argmin}} \left[\frac{1}{n_a} \sum_{i:\, A_i = a} \{Y_i - \hat{\mu}_a^o(X_i) - d(X_i)\}^2 + \mathcal{P}_a(d) \right]$$

n_a = number of RCT subjects with $A_i = a$

$\mathcal{P}_a(d)$ = penalty on $d(x)$ (e.g., ridge, spline, GP)

➡ **Penalty shrinks discrepancy toward 0, preserving borrowing from OS.**

Takeaway: R-OSCAR **borrows** the stable OS outcome model $\hat{\mu}_a^o(x)$, then learns a **penalized discrepancy** $\hat{\delta}_a(x)$ on RCT data to yield $\hat{\mu}_a^{cal}(x) = \hat{\mu}_a^o(x) + \hat{\delta}_a(x)$ —aligning with the RCT **without relearning** the full RCT outcome model.

Covariate mismatch

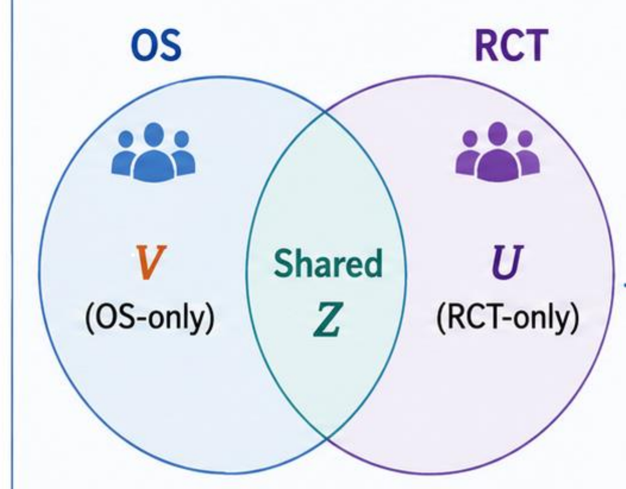

Goal: estimate CATE in the RCT population

1 RACER

RCT → Arm-specific outcome models → Pseudo-outcome regression → RCT CATE

RCT only

2 SR-OSCAR

OS → OS model on Z; RCT → RCT calibration using (U, Z) → Pseudo-outcome regression → RCT CATE

borrows through shared Z

3 MR-OSCAR

OS → Impute V from Z → OS model on (Z, V); RCT → RCT calibration using $(U, Z, \hat{V})$ → Pseudo-outcome regression → RCT CATE

uses predicted OS-only covariates

Why borrowing from OS can improve CATE estimation

Step 1: CATE error is driven by outcome-model accuracy

$$\Delta_2^2(\hat{\tau}, \tau^r) \lesssim \Delta_2^2(\mathcal{F}, \tau^r) + \sum_{a \in \{-1,1\}} \Delta_2^2(\hat{\mu}_a, \mu_a^r)$$

CATE estimation error is bounded by approximation error of the CATE model class plus the sum of arm-wise outcome-model errors—better arm-wise outcome models lead directly to better CATE estimation.

Step 2: How borrowing improves outcome-model estimation

RACER (RCT only)

$$\Delta_2^2(\hat{\mu}_a^r, \mu_a^r) \lesssim \frac{c(\mathcal{M}_a^r)}{(n_a^o)^{n_a^o}}$$

R-OSCAR (borrow + calibrate)

$$\Delta_2^2(\hat{\mu}_a^{o,\dagger}, \mu_a^r) \lesssim \frac{c(\mathcal{M}_a^o)}{(n_a^o)^{n_a^o}} + \frac{c(\mathcal{D}_a^\dagger)}{(n_a^r)^{n_a^r}}$$

Borrowing helps when, for all a, $C_1 \frac{c(\mathcal{M}_a^o)}{(n_a^o)^{n_a^o}} + C_2 \frac{c(\mathcal{D}_a^\dagger)}{(n_a^r)^{n_a^r}} < C_3 \frac{c(\mathcal{M}_a^r)}{(n_a^r)^{n_a^a}}$

Large OS sample + simpler discrepancy correction can beat direct RCT-only estimation.

Step 3: Sparse linear special case

RACER (RCT only)

$$\Delta_2^2(\hat{\mu}_a^r, \mu_a^r) \approx \frac{p \log(p)}{n_a^o}$$

R-OSCAR (borrow + calibrate)

$$\Delta_2^2(\hat{\mu}_a^{o,\dagger}, \mu_a^r) \approx \frac{p \log(p)}{n_a^o} + \frac{s \log(p)}{n_a^r}$$

If $n_a^o \gtrsim \frac{p}{p-s} n_a^r$, then **R-OSCAR** improves over **RACER**.

OS learns the full high-dimensional signal; the RCT only learns a sparse correction.

MR-OSCAR improves over **SR-OSCAR** when Z predicts V well and the imputation cost is smaller than the error from omitting V.

A.

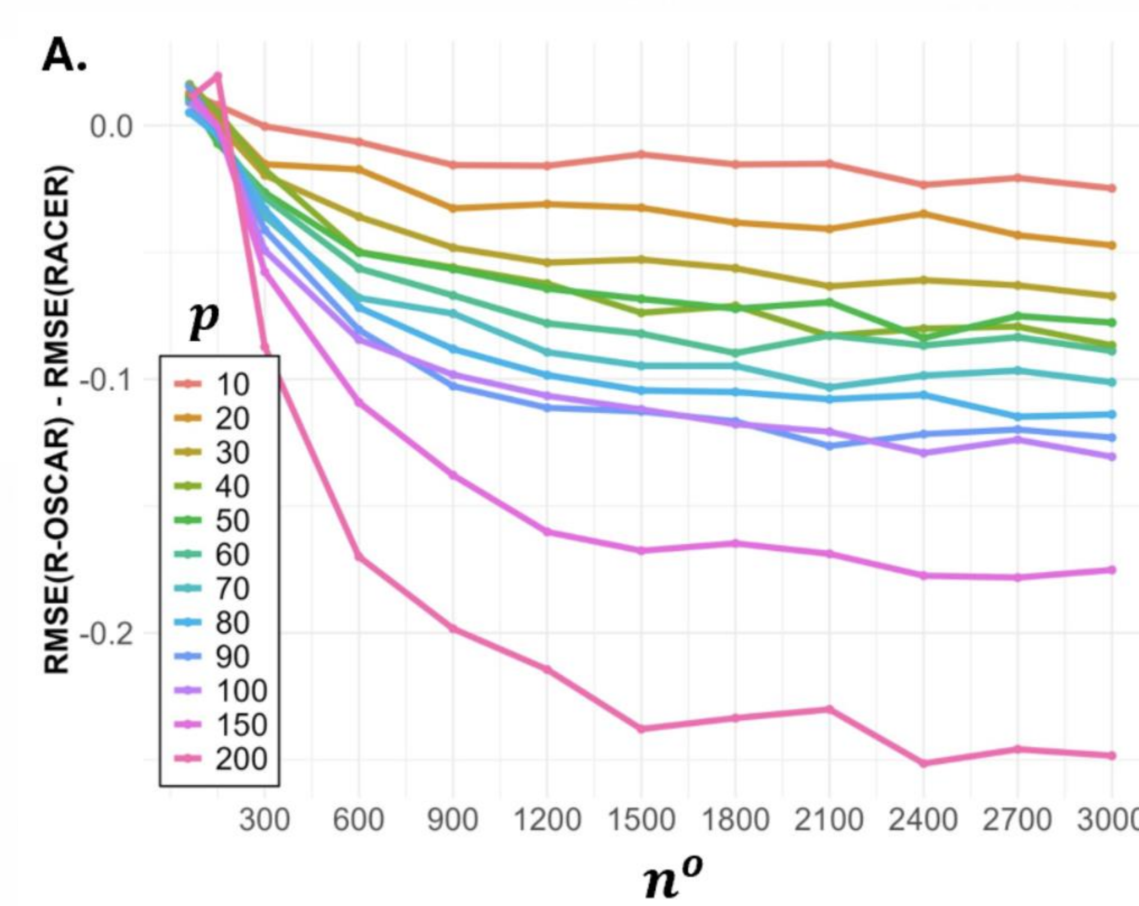

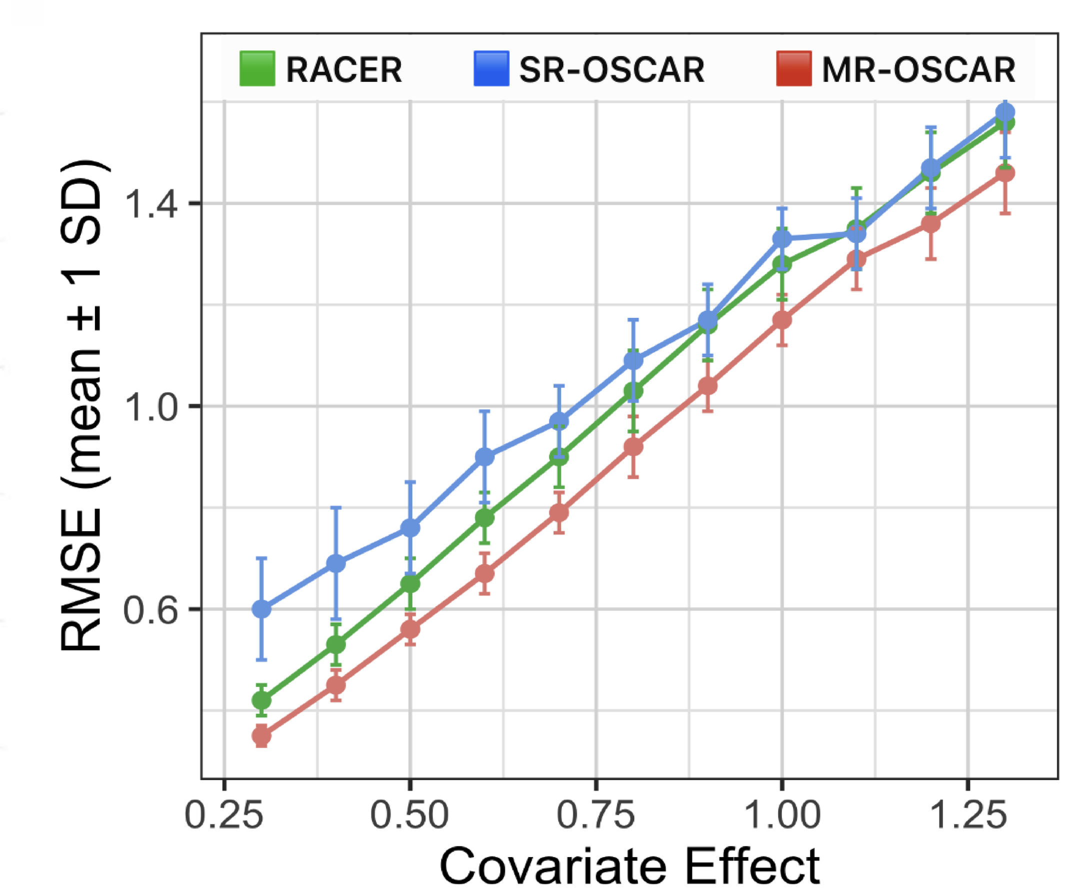

Main Takeaway: Better outcome modeling leads to better CATE estimation; borrowing helps when observational signal is strong, calibration is simple, and mismatch is controlled.

Contact

Samhita Pal: samhita.pal@vumc.org

Amir Asiaee: amir.asiaeetaheri@vumc.org

Jared D Huling: huling@umn.edu

ME-2023C1-32148

PATIENT-CENTERED OUTCOMES RESEARCH INSTITUTE