

Demystifying TMLE: A Practical Approach for Clinical Data Scientists

Joshua Enxing

Phastar

joshua.enxing@phastar.com

March 23rd, 2026
PHUSE US Connect




Outline

- 1 Clinical trial questions vs statistical models
- 2 Estimands and common analysis targets
- 3 Collapsibility and non-collapsibility
- 4 Post-treatment covariates and censoring
- 5 What TMLE is (definition and intuition)
- 6 Key ingredients: EIF and targeting (high level)
- 7 Simple randomized trial example
- 8 Complex longitudinal example
- 9 Diagnostics and sensitivity analyses
- 10 Key takeaways

Clinical Trials Ask Intervention-Based Questions

- Typical scientific question:
“What would happen to outcomes if all patients received Treatment A instead of Treatment B?”
- This is a **population-level** question
- The answer should not depend on arbitrary modeling choices



What Is an Estimand?

Estimand:

- A precisely defined quantity that answers the scientific question
- Examples:
 - Average Treatment Effect (ATE)
 - Risk difference at a fixed time
 - Mean outcome under a treatment strategy

$$\text{ATE} = E[Y^{(1)}] - E[Y^{(0)}]$$

- $Y^{(a)}$: potential outcome under treatment a
- This is a **marginal** (population-average) quantity


Common Practice: Covariate-Adjusted Regression

$$P(Y = 1 \mid A, W) = \beta_0 + \beta_1 A + \beta^\top W$$

- A : treatment
- W : baseline covariates
- Reported effect: $\exp(\beta_1)$

Key question: What estimand does $\exp(\beta_1)$ correspond to?

Conditional vs Marginal Effects



- $\exp(\beta_1)$ is a **conditional odds ratio**
- It compares treated vs untreated patients:
 - holding W fixed
 - within covariate strata
- The clinical estimand is usually **marginal**

Definition: Collapsibility

An effect measure is **collapsible** if adjusting for a covariate that is *not* a *confounder* does not change the effect.

Formally, an effect measure θ is collapsible over Z if:

$$\theta(Y, A) = \sum_z \theta(Y, A \mid Z = z) P(Z = z)$$


- Mean difference: collapsible
- Risk difference: collapsible
- Odds ratio: **not collapsible**
- Hazard ratio: **not collapsible**

Non-Collapsibility: Why It Matters

- Odds ratios and hazard ratios are **non-collapsible**
- Even if:
 - treatment is randomized
 - covariates are purely prognostic
 - model is correctly specified
- Adjusting for W changes the estimand

Post-Treatment Covariates

- Many trials measure post-baseline biomarkers
- These variables are affected by treatment
- Conditioning on them can:
 - block causal pathways
 - change the estimand



Informative Censoring

- Outcomes may be missing or censored
- If censoring depends on covariates or treatment:
 - complete-case analysis is biased
- Conditioning on being observed induces selection bias

What Is TMLE?

Targeted Maximum Likelihood Estimation (TMLE):

- A framework for estimating causal estimands
- Targets the estimand directly
- Uses covariates for efficiency
- Provides valid inference under flexible modeling

What TMLE Targets

$$\Psi = E[E(Y | A = 1, W) - E(Y | A = 0, W)]$$

- Marginal estimand
- Averaged over the observed population
- Independent of modeling parameterization


Efficient Influence Function (EIF)

EIF:

- A function describing how each observation influences the estimand
- Determines optimal asymptotic variance
- Central to TMLE estimation and inference

$$D^*(O) = \frac{I(A = 1)}{g(1 | W)}(Y - Q(1, W)) - \frac{I(A = 0)}{g(0 | W)}(Y - Q(0, W)) \\ + Q(1, W) - Q(0, W) - \Psi$$


TMLE Algorithm (High Level)




- 1 Estimate outcome regression $Q(A, W)$
- 2 Estimate treatment mechanism $g(A | W)$
- 3 Construct the clever covariate
- 4 Update Q via targeting
- 5 Plug in to estimate Ψ

Example 1: Randomized Trial with Baseline Adjustment

- Binary outcome, randomized treatment
- Logistic regression $\rightarrow$ conditional OR
- TMLE $\rightarrow$ marginal mean difference
- Covariates improve precision, not estimand



Example 2: Longitudinal Strategy



- Treatment initiation and adaptation
- Post-treatment biomarkers
- Informative censoring
- Estimand: mean outcome under a dynamic regime

Diagnostics You Should Always Report

- Overlap / positivity
- Clever covariate and EIF stability
- Transparency of machine learning choices
- Sensitivity across truncation and methods

Key Takeaways

- Covariate adjustment can change the estimand
- Non-collapsibility explains why
- TMLE targets marginal, policy-relevant effects
- Diagnostics are essential for credibility

References

- 1 van der Laan, M.J. and Rose, S. *Targeted Learning: Causal Inference for Observational and Experimental Data*. Springer Series in Statistics, New York, NY: Springer; 2011.
- 2 Talbot, D., Diop, A. et al. Guidelines and best practices for the use of targeted maximum likelihood and machine learning when estimating causal effects of exposures on time-to-event outcomes. *Statistics in Medicine*, **44**: e70034. <https://doi.org/10.1002/sim.70034>
- 3 Havlir, D.V., Balzer, L.B., Charlebois, E.D., et al. HIV testing and treatment with the use of a community health approach in rural Africa. *The New England Journal of Medicine*. 2019; **381**(3):219–229.
- 4 Balzer, L.B., Ayieko, J., Kwarisiima, D., et al. Far from MCAR: obtaining population-level estimates of HIV viral suppression. *Epidemiology*. 2020;31(5):620–627. doi:10.1097/EDE.0000000000001215.
- 5 Hejazi, N.S., van der Laan, M.J., Janes, H.E., Gilbert, P.B., and Benkeser, D.C. Efficient nonparametric inference on the effects of stochastic interventions under two-phase sampling, with applications to vaccine efficacy trials. *Biometrics*. 2021;77(4):1241–1253. doi:10.1111/biom.13375.

Thank you for listening! Questions?