



Modeling Covariate Transition for Efficient Estimation of Longitudinal Treatment Effects in Randomized Experiments

Naoki Chihara¹ Tatsushi Oka² Yasuko Matsubara¹ Yasushi Sakurai¹ Shota Yasui³¹SANKEN, The University of Osaka ²Keio University ³CyberAgent, Inc.

Motivation

- Regression-adjustment** methods are useful for variance reduction in randomized experiments, but standard methods mainly use pre-treatment covariates and estimate a single average effect
- It is important to capture insights into the temporal evolution of treatment effects (e.g., immediate, delayed, or persistent effects)
- Post-treatment histories are key covariates, but naive conditioning may change the causal estimand

Problem Setup

Our work focuses on longitudinal randomized experiments under static regimes

$$\{(\bar{X}_{i,T}, \bar{W}_{i,T}, \bar{Y}_{i,T})\}_{i=1}^n \stackrel{\text{i.i.d.}}{\sim} P_{\bar{Z}_T}$$

- $Y_t \in \mathcal{Y}_t \subseteq \mathbb{R}$: Outcome at time t
- $W_t \in \mathcal{W}_t$: Treatment assignment at time t
- $X_t \in \mathcal{X}_t \subseteq \mathbb{R}^d$: Covariates at time t

Goal of this work

Estimate the time-varying effects $\{\text{ATE}(t)\}_{t=1}^T$ not just a single cumulative effect, where

$$\text{ATE}(t) = \mu_{\bar{w}_t}(t) - \mu_{\bar{w}'_t}(t), \quad \mu_{\bar{w}_t}(t) = \mathbb{E}[Y_t(\bar{w}_t)].$$

The empirical estimator is unbiased and consistent,

$$\hat{\mu}_{\bar{w}_t}^{\text{emp}}(t) := \frac{1}{n_{\bar{w}_t}} \sum_{i: \bar{W}_{i,t} = \bar{w}_t} Y_{i,t},$$

- Since predictive histories can improve its precision, we utilize the following mean regression

$$m_{\bar{w}_t}^{(t)}(\bar{h}_t) = \mathbb{E}[Y_t(\bar{w}_t) | \bar{H}_t(\bar{w}_{t-1}) = \bar{h}_t]$$

- Challenges: it receives post-treatment histories

Methodology

To represent the dynamic trajectories in the historical data, we utilize the **transition kernels**

Measure Factorization via Transition Kernels

$$P_{\bar{H}_t}^{\bar{w}_{t-1}}(d\bar{h}_t) = P_{H_1}(dh_1) \prod_{\tau=1}^{t-1} p_{\bar{w}_\tau}^{(\tau)}(dh_{\tau+1} | \bar{h}_\tau).$$

Forward-Integrated Regression

$$\Gamma_{\bar{w}_t}^{(\tau)}(\bar{h}_\tau) = \int_{\mathcal{H}_{\tau+1}} \Gamma_{\bar{w}_t}^{(\tau+1)}(\bar{h}_{\tau+1}) p_{\bar{w}_\tau}^{(\tau)}(dh_{\tau+1} | \bar{h}_\tau), \quad \Gamma_{\bar{w}_t}^{(t)}(\bar{h}_t) = m_{\bar{w}_t}^{(t)}(\bar{h}_t).$$

Our estimator can capture temporal dynamics while avoiding post-treatment bias thanks to transition kernels

Dynamic Regression-Adjusted Estimator

$$\begin{aligned} \hat{\mu}_{\bar{w}_t}^{\text{dy-adj}}(t) &= \frac{1}{n_{\bar{w}_t}} \sum_{i: \bar{W}_{i,t} = \bar{w}_t} \left(Y_{i,t} - \hat{\Gamma}_{\bar{w}_t}^{(t)}(\bar{H}_{i,t}) \right) \\ &\quad + \frac{1}{n} \sum_{i=1}^n \hat{A}_{\bar{w}_t}^{(t)}(\bar{H}_{i,t}) + \frac{1}{n} \sum_{i=1}^n \hat{\Gamma}_{\bar{w}_t}^{(1)}(X_{i,1}), \end{aligned}$$

where $\hat{A}_{\bar{w}_t}^{(t)}$ is correction term for transition innovations

Theory

Summary of Theoretical Results

Under suitable assumptions,


- Neyman orthogonality with cross-fitting gives valid statistical inference
- Our estimator satisfies $\sqrt{n}(\hat{\theta}_t - \theta_t) \rightsquigarrow \mathcal{N}(0, \Sigma_t)$, and attains the semiparametric efficiency bound Σ_t
- Oracle dynamic adjustment is asymptotically no worse than the empirical estimator

Simulation Study

Synthetic datasets are generated according to

$$\begin{aligned} Y_{i,t} &= g(X_{i,t}, Y_{i,t-1}) + \tau(t)W_{i,t} + \varepsilon_{i,t}, \\ dX_{i,t} &= Q \cdot f(Q^\top X_{i,t}, Y_{i,t-1}, W_{i,t})dt + \sigma_X dB_t, \end{aligned}$$

- (1) nonlinearity, (2) irrelevant covariates, (3) influences from past treatments, and (4) time-varying effects




- Our estimator achieves improvements in both RMSE and CI length compared to empirical baseline, while maintaining the nominal coverage probability
- While linear adjustment moderately enhances precision, a nonlinear one achieves further improvements

Application to Real-World Datasets

Our framework is applied to A/B test data from a large-scale streaming platform to evaluate its practicality

- 10-day A/B test for up-next content recommendations
- Outcome: daily viewing time of a target series
- Control group: user interaction based recommendations
- Treatment group: recommendations based on content-level similarity from item features



- Standard errors are reduced by about 0.4-20.4%