Nonparametric doubly-robust inference for the mean outcome under a longitudinal treatment decision rule

David Benkeser
Emory Univ.

Marco Carone
Univ. of Washington

Peter Gilbert
Fred Hutch.

Georgia Statistics Day

10/26/2018

Motivation for considering time-varying interventions

In many clinical contexts, the treatment of interest is administered in phases over time.

- antihypertensive drug therapy administered daily;
- biphosphonate drug therapy administered weekly;
- injection of antiretroviral suspension administered every month;
- immunosuppressant infusion therapy administered every two months.

The observed data are often of the form

$$\underbrace{L_0 \longrightarrow A_0}_{\text{time } 0} \longrightarrow \underbrace{L_1 \rightarrow A_1}_{\text{time } 1} \longrightarrow \cdots \longrightarrow \underbrace{L_K \longrightarrow A_K}_{\text{time } K} \longrightarrow Y\,,$$

where we have defined components

$$\begin{aligned} L_k &= \text{covariates recorded at time } k; \\ A_k &= \text{treatment assignment at time } k; \\ Y &= \text{outcome recorded at the end of the study.} \end{aligned}$$

Motivation for considering time-varying interventions

We can consider the counterfactual outcome $Y(a_0, a_1, \ldots, a_K)$ defined by enforcing treatment assignment $(A_0, A_1, \ldots, A_K) = (a_0, a_1, \ldots, a_K)$.

This allows to define causal contrasts that address the scientific question of interest.

(Chapters 24-26 of van der Laan & Rose, 2011; Chapter 4 of of van der Laan & Rose, 2018; Chapter 19 of Hernán & Robins, 2018)

Weekly alendronate therapy for osteoporosis and one-year incidence of hip fracture:

L_k = covariates recorded at week k
(e.g., sex, age, dexascan values, thyroid hormone levels, side effects, fracture status);
A_k = indicator that alendronate was taken at week k;
Y = indicator that hip fracture occurred within one year.

We may be interested in the average effect

$$E[Y(1, 1, \ldots, 1)] - E[Y(0, 0, \ldots, 0)]$$

of year-long weekly alendronate therapy on one-year risk of hip fracture versus no alendronate therapy, or other contrasts defined by values of $(a_0, a_1, \ldots, a_{52})$.

Motivation for considering time-varying interventions

Even when the treatment is administered at a single time-point, it is often the case that the data are incompletely recorded in the follow-up period.

- missing data: patient did not show up to a scheduled clinic visit;
- loss to follow-up: patient moved out of the country and dropped out of the study.

It would be natural then to consider a counterfactual outcome defined by enforcing

1. the administration of a particular treatment (baseline only or time-varying);
2. complete follow-up and complete recording of data (time-varying).

What would the outcome have been had:

> the patient taken an experimental treatment regime, the follow-up been complete, and all data been completely recorded?

> the patient taken a control treatment regime, the follow-up been complete, and all data been completely recorded?

Motivation for considering time-varying interventions

For example, if treatment is only administered at baseline, we could set:

L_k = covariates recorded at time k;

A_0 = treatment assignment at time 0 (i.e., at baseline);

A_k = indicator that, at time k, patient has not yet been lost to follow-up and all measurements on this patient are complete;

Y = outcome recorded at the end of the study.

We might then be interested in

$$ATE = E[Y(1,1,1,\ldots,1)] - E[Y(0,1,1,\ldots,1)] \,.$$

Motivation for considering time-varying interventions

If treatment is administered over time, we could instead set:

L_k = covariates recorded at time k;

$A_{k,1}$ = indicator that, at time k, patient has not yet been lost to follow-up and all measurements on this patient are complete;

$A_{k,2}$ = indicator of treatment assignment at time k;

Y = outcome recorded at the end of the study.

and let $Y(\underbrace{(a_{0,1}, a_{0,2})}_{a_0}, \underbrace{(a_{1,1}, a_{1,2})}_{a_1}, \ldots, \underbrace{(a_{K,1}, a_{K,2})}_{a_K})$ be the counterfactual defined by

$$(A_0, A_1, \ldots, A_K) = (a_0, a_1, \ldots, a_K) \ ,$$

where we write $A_k := (A_{k,1}, A_{k,2})$.

We might then be interested in

$$E[Y((1,1),(1,1),\ldots,(1,1))] - E[Y((1,0),(1,0),\ldots,(1,0))] \ .$$

Dynamic treatment rules

Counterfactuals defined by fixed treatment profiles are often neither particularly clinically interesting nor supported by data.

Treatment decisions are usually dynamic and incorporate real-time patient information.

Example: mercaptopurine in IBD patients

- static intervention: 'always treat' versus 'never treat'
- if patient develops signs of liver damage, therapy is usually stopped
- liver function is a time-varying confounder between treatment status and survival
- if poor liver function is a contraindication for therapy, it may not be possible to observe treatment adherence among patients with recent liver failure
- static intervention is unrealistic and not identifiable
- dynamic intervention: 'treat while liver function permits it' versus 'never treat'

$$d(t) = \begin{cases} 1 & : \text{ if recent liver function is adequate} \\ 0 & : \text{ otherwise} \end{cases} .$$

Dynamic treatment rules

Counterfactuals can be naturally defined in terms of **dynamic treatment rules encoding treatment decisions that possibly depend on current and past patient info**.

In the mercaptopurine example, we may want to learn about the average effect

$$ATE(d, d_0) := E[Y(d)] - E[Y(d_0)]$$

of rule d enforcing treatment whenever liver function permits it and rule d_0 enforcing no mercaptopurine use.

Our goal is to contrast the mean outcome under various sequences of interventions occurring over time. To simplify notation, we focus on static treatment profile $(a_0, a_1, \ldots, a_K) = (1, 1, \ldots, 1)$, but methods easily extend to dynamic treatment rules.

Identification of average treatment effects

Our goal is to contrast the mean outcome under various sequences of interventions occurring over time. To simplify notation, we focus on static treatment profile $(a_0, a_1, \ldots, a_K) = (1, 1, \ldots, 1)$, but methods easily extend to dynamic treatment rules.

(Sequentially) randomized trial

We can imagine conducting a trial in which, at each of these time-points, individuals are randomized to one of the possible interventions.

In this case, at each time-point, the intervention assignment is independent of the possible counterfactual outcomes.

$$Y(1, 1, \ldots, 1) \perp A_0 \ , \quad Y(1, 1, \ldots, 1) \perp A_1 \ , \quad \ldots \ , \quad Y(1, 1, \ldots, 1) \perp A_K \ .$$

Identification of average treatment effects

Our goal is to contrast the mean outcome under various sequences of interventions occurring over time. To simplify notation, we focus on static treatment profile $(a_0, a_1, \ldots, a_K) = (1, 1, \ldots, 1)$, but methods easily extend to dynamic treatment rules.

Observational study

In an observational study, there are often factors that influence both the intervention assignment mechanisms and the counterfactual outcome distribution.

Examples of time-varying confounding:

- a patient may discontinue chemotherapy because they have ceased to respond, which may itself be a marker of disease progression;
- a patient may have ceased smoking because they developed respiratory symptoms, which may be a sign of lung cancer.

Identification of average treatment effects

Our goal is to contrast the mean outcome under various sequences of interventions occurring over time. To simplify notation, we focus on static treatment profile $(a_0, a_1, \ldots, a_K) = (1, 1, \ldots, 1)$, but methods easily extend to dynamic treatment rules.

Observational study

The vector of time-varying covariates $(L_0, L_1, \ldots, L_K)$ can be used to deconfound the relationship between Y and $(A_0, A_1, \ldots, A_K)$ provided

$$Y(1,1,\ldots,1) \perp A_0 \mid L_0 \ , \quad Y(1,1,\ldots,1) \perp A_1 \mid \overline{L}_1, A_0 = 1 \ , \quad \ldots$$
$$Y(1,1,\ldots,1) \perp A_K \mid \overline{L}_K, \overline{A}_{K-1} = 1_K \ ,$$

where the symbol 1_j is used to denote a vector $(1, 1, \ldots, 1)$ of length j.

In other words, at each time-point, intervention assignment is randomized *within each stratum* defined by recorded patient history up to that point, among patients who have received the intervention of interest so far.

This is referred to as the **sequential randomization** (or **exchangeability**) **condition**.

Identification of average treatment effects

Our goal is to infer what the mean outcome would be in the target population under the multi time-point intervention of interest.

We must be able to observe the intervention of interest for each different "type" of individual (as defined by recorded covariates) from this population:

- $g_0(\ell_0) = P(A_0 = 1 \mid L_0 = \ell_0) > 0$ for each possible ℓ_0;
- $g_1(\ell_1) = P(A_1 = 1 \mid \bar{L}_1 = \bar{\ell}_1, A_0 = 1) > 0$ for each possible $\bar{\ell}_1$;
- $\ldots$
- $g_K(\ell_K) = P(A_K = 1 \mid \bar{L}_K = \bar{\ell}_K, \bar{A}_K = 1_K) > 0$ for each possible $\bar{\ell}_K$.

This is referred to as the **positivity condition**.

Identification of average treatment effects

Under these conditions, the mean counterfactual $E[Y(1,1,\ldots,1)]$ equals the multi time-point **G-computation formula** (Robins, 1986).

$$E\Big[\, E\Big[\, E\Big[\ldots \Big[\, E\Big[\, \underbrace{E\Big(Y \,\Big|\, \overline{A}_K = 1, \overline{L}_K\Big)}_{\bar{Q}_{K+1}(\overline{L}_K)} \,\Big|\, \overline{A}_{K-1} = 1, \overline{L}_{K-1}\Big] \ldots \Big]\Big|\, \overline{L}_1, A_0 = 1\Big]\Big|\, L_0\Big]\Big] .$$

$\bar{Q}_K(\overline{L}_{K-1})$

$\bar{Q}_2(\overline{L}_1)$

$\bar{Q}_1(L_0)$

$\bar{Q}_0$

where, for any k, we write $\overline{A}_k := (A_0, A_1, \ldots, A_k)$ and $\overline{L}_k := (L_0, L_1, \ldots, L_k)$.

The (mathematically equivalent) **IPTW identification formula** is given by

$$E[Y(1,1,\ldots,1)] = E\left[\left\{\frac{A_0 A_1 \ldots A_K}{\bar{g}_K(\overline{L}_K)}\right\} Y\right] ,$$

where, for any k, we write $\bar{g}_k = \prod_{m=1}^{k} g_m$.

Identification of average treatment effects

The equivalence between the IPTW and G-computation identification formulas can be established through repeated uses of the law of total expectation.

$$
\begin{aligned}
E\left[\left\{\frac{A_0 A_1 \ldots A_K}{g_0(L_0) g_1(\overline{L}_1) \ldots g_K(\overline{L}_K)}\right\} Y\right] &= E\left[E\left[\left\{\frac{A_0 A_1 \ldots A_K}{g_0(L_0) g_1(\overline{L}_1) \ldots g_K(\overline{L}_K)}\right\} Y \,\middle|\, \overline{L}_K, \overline{A}_K\right]\right] \\
&= E\left[\left\{\frac{A_0 A_1 \ldots A_K}{g_0(L_0) g_1(\overline{L}_1) \ldots g_K(\overline{L}_K)}\right\} \bar{Q}_{K+1}(\overline{L}_K)\right] \\
&= E\left[\left\{\frac{A_0 A_1 \ldots A_{K-1}}{g_0(L_0) g_1(\overline{L}_1) \ldots g_{K-1}(\overline{L}_{K-1})}\right\} \frac{\bar{Q}_{K+1}(\overline{L}_K)}{g_K(\overline{L}_K)} E\left(A_K \,\middle|\, \overline{L}_K, \overline{A}_{K-1}\right)\right] \\
&= E\left[\left\{\frac{A_0 A_1 \ldots A_{K-1}}{g_0(L_0) g_1(\overline{L}_1) \ldots g_{K-1}(\overline{L}_{K-1})}\right\} E\left[\bar{Q}_{K+1}(\overline{L}_K) \,\middle|\, \overline{L}_{K-1}, \overline{A}_{K-1} = 1_{K-1}\right]\right] \\
&= E\left[\left\{\frac{A_0 A_1 \ldots A_{K-1}}{g_0(L_0) g_1(\overline{L}_1) \ldots g_{K-1}(\overline{L}_{K-1})}\right\} \bar{Q}_K(\overline{L}_{K-1})\right] = \ldots
\end{aligned}
$$

Estimation procedures

The G-computation and IPTW formulas suggest natural estimation strategies.

- For G-computation, sequentially estimate outcome regressions.
- For IPTW, estimate propensity for treatment at each time.

There are also frameworks for combining the two approaches including

- **augmented inverse probability of treatment weighting (AIPTW)**
- **targeted minimum loss-based estimation (TMLE)**.

There are several benefits to considering these more complex frameworks.

- Nonparametric efficient estimation if **outcome regressions and propensity scores are both consistent** for their true respective counterparts.
- Consistent estimation if **either outcome regressions or propensity scores are consistent** for their true respective counterparts.
- The latter property is known as **double-robustness**.

Inferential procedures

What about inference? We would like to formally compare (e.g., test) differences in average outcomes under different treatment strategies.

If outcome regressions and propensity scores can be consistently estimated via parametric regressions, inference may be facilitated via the nonparametric bootstrap.

- In practice, correctly specifying a single parametric model is challenging.
- Here, we require at least K correct parametric regression models!

Modern statistical learning methods (e.g., machine learning) offer an alternative to classic parametric approaches.

- Better chance of getting either outcome regressions or propensity score correct.
- **BUT** inference is only valid if both are correct.
- When one is incorrect, naive confidence intervals have poor coverage.
- Nonparametric bootstrap is not generally valid when using these methods.

Inferential procedures

Why the poor performance of standard inference?

- When one of the outcome regression or propensity score is inconsistent, the bias of the estimate of the counterfactual shrinks slower than $n^{-1/2}$.
- Standard Wald-style confidence intervals and hypothesis tests are based on standard error estimates that shrinks at rate $n^{-1/2}$,

$$\hat{\sigma} = \frac{\widehat{\text{Var}}\{\hat{E}[Y(1,\ldots,1)]\}}{n^{1/2}} .$$

- Intervals shrink with $n^{-1/2}$, but center around the truth at a slower rate $\Rightarrow$ **asymptotic coverage probability of 0% and type-I error rate of 1!**

Double-robust inference

To correct for this, we require a better understanding of how inconsistent estimation of a nuisance parameter generates bias in the estimate of the target parameter.

- Requires characterization of the second-order remainder of the von Mises expansion of target parameter.

It is possible to use the **TMLE framework** to construct an estimator that

1. is efficient when both outcome regression and propensity scores are consistent;
2. is consistent when at least one is consistent;
3. when suitably normalized, tends to a mean-zero normal distribution with variance we can consistently estimate, when at least one is consistent.

It does not appear possible to adapt the AIPTW estimator for this purpose.

Details for a single timepoint intervention are provided in Benkeser, Carone, van der Laan & Gilbert (2017) and in the `R` package `drtmle` (available on CRAN).

Summary: estimation and inference

Properties of estimation procedures outlined

		$\bar{Q} + \bar{g}$		$\bar{Q} + \bar{g}$		$\bar{Q} + \bar{g}$	
	difficulty	**target**	**ci**	**target**	**ci**	**target**	**ci**
IPTW	+			✓		✓	
G-COMP	++	✓				✓	
AIPTW	+++	✓		✓		✓	✓
TMLE	++++	✓	✓	✓	✓	✓	✓

$\bar{Q} + g$: outcome regressions estimated well but not propensity scores
$\bar{Q} + g$: propensity scores estimated well but not outcome regressions
$\bar{Q} + g$: outcome regressions and propensity scores estimated well

target : does the estimator hit the right target?
ci : is valid inference possible and readily available, even when flexible learning regression strategies are used?

Key points

- Methods for time-varying interventions are extremely versatile, and can be used to tackle loss to follow-up and missing data.
- Dynamic treatment rules may better reflect realistic interventions and prevent positivity violations.
- Doubly-robust estimators should be preferred as they confer efficiency, additional robustness and the ability to use flexible regression estimators.
- Naive inference with inconsistent regression estimates can be disastrous. Additional steps are needed to ensure valid inference.

References and additional reading

References:

Bang H, Robins JM (2005). Doubly robust estimation in missing data and causal inference models. *Biometrics*; 61(4)962-973. doi: 10.1111/j.1541-0420.2005.00377.x/

Benkeser, D, Carone, M, van der Laan, MJ & Gilbert, P (2017). Doubly-robust nonparametric inference on the average treatment effect. *Biometrika*; 104(4)863-880. doi: 10.1093/biomet/asx053.

Hernán MA, Robins JM (2018). *Causal Inference*. Chapman & Hall/CRC. *Forthcoming – draft available online*.

Robins JM (1986). A new approach to causal inference in mortality studies with a sustained exposure period – application to control of the healthy worker survivor effect. *Mathematical Modelling*, 9(7)1393-1512. doi: 10.1016/0270-0255(86)90088-6.

van der Laan MJ, Rose S (2011). *Targeted Learning: Causal Inference for Observation and Experimental Data*. Springer New York. 10.1007/978-1-4419-9782-1

van der Laan MJ, Gruber S (2012). Targeted minimum loss based estimation of causal effects of multiple time point interventions. *International Journal of Biostatistics*; 8(1). doi: 10.1515/1557-4679.1370

van der Laan MJ, Rose S (2018). *Targeted Learning in Data Science: Causal Inference for Complex Longitudinal Studies*. Springer New York. 10.1007/978-3-319-65304-4