

Statistical Learning in Mediation Analysis

Chapter 6: Direct and indirect effects for stochastic interventions

David Benkeser
Emory University

Iván Díaz
New York University

Marco Carone
University of Washington

MODULE 14

**Summer Institute in Statistics for
Clinical and Epidemiological Research**

July 2024

Contents of this chapter

Total effects:

- 1 What are stochastic interventions
- 2 Modified treatment policies and incremental propensity score interventions
- 3 Identification of effects of stochastic interventions
- 4 R packages and examples

Mediation effects:

- 1 Defining mediation effects for stochastic interventions
- 2 When and how are these effects identified
- 3 Estimation based on the G-computation and IPW formulas
- 4 Doubly-robust estimation
- 5 R packages and examples

What are stochastic interventions?

- Consider a simple data structure $O = (W, A, Y)$
- Recall the counterfactual definitions $Y(1)$ and $Y(0)$
- Obtained by considering hypothetical worlds where A is set to $A = 1$ and $A = 0$ with probability one.

We will alternatively consider interventions where A is set to some random variable A_δ with probability one.

Examples of useful stochastic interventions

Modified treatment policies (MTP)

We focus MTPs, which are one type of intervention where the post-intervention exposure is a modification of the actual exposure $A_\delta = d(A, W)$. Examples:

- Let A denote a measure of air quality given by particulate matter $\text{PM}_{2.5}$. One could be interested in an intervention that would reduce $\text{PM}_{2.5}$ by 10%:

$$d(A, W) = (1 - \delta)A,$$

where $\delta = 0.1$.

- Let A denote self-reported physical activity as measured by weekly minutes. One could be interested in an intervention that increases physical activity for people for whom it is feasible:

$$d(A, W) = \begin{cases} A + \delta & \text{if } A + \delta < u(W) \\ A & \text{if } A + \delta \geq u(W) \end{cases},$$

where $u(W)$ is the upper bound of physical activity for someone with covariates W (age, health status, etc.).

Examples of useful stochastic interventions

Incremental propensity score interventions (IPSI):

Let ϵ denote a draw from $U(0, 1)$. The MTP is given by

$$d(A, W) = \begin{cases} A & \text{if } \epsilon < \delta \\ 0 & \text{otherwise,} \end{cases}$$

where $0 < \delta < 1$ is a user given value.

Note that:

$$g_{A,\delta}(1 | w) = \delta g_A(1 | w)$$

Note: Kennedy, 2019 proposed interventions where instead δ is an odds ratio.

Definition of counterfactuals and causal effects

Define the counterfactual variable $Y(A_\delta)$ as the variable that **would have been observed in a hypothetical world where $A = A_\delta$** .

- We can contrast the expectation of $Y(A_\delta)$ with the expectation of Y to obtain a causal effect:

$$E[Y(A_\delta) - Y] \quad \text{or} \quad E[Y(A_\delta)] - E[Y]$$

- We know how to estimate $E(Y)$ well (using the empirical mean)
- In what follows **we focus on identifying and estimating $E[Y(A_\delta)]$**

Identification of the mean outcome under a stochastic intervention

We need the usual two assumptions:

- **Positivity:** if $g_A(a | w) = 0$ then $g_{A,\delta}(a | w) = 0$.
- **Randomization:** $A \perp Y(a) | W$ for all a

We have that

$$E[Y(A_\delta)] = E[\bar{Q}_Y(A_\delta, W)],$$

where, as before, $\bar{Q}_Y(a, w) := E(Y | A = a, W = w)$.

Positivity assumption

Positivity assumption for IPSIs:

- Recall that

$$g_{A,\delta}(1 | w) = \delta g_A(1 | w)$$

- $g_A(1 | w) = 0$ implies $g_{A,\delta}(1 | w) = 0$.

Positivity assumption for MTPs:

- Let's look at one of the examples:

$$d(A, W) = \begin{cases} A + \delta & \text{if } A + \delta < u(W) \\ A & \text{if } A + \delta \geq u(W) \end{cases},$$

- Assume $P(A < u(W) | W) = 1$
- Then $g_A(a | w) = 0$ implies $g_{a,\delta}(a | w) = 0$

When are stochastic interventions useful?

- To define **meaningful effects for non-binary exposures**:
 - In some applications (e.g., physical activity) **it may make little sense to work with counterfactuals that set $A = a$**
 - Even if defining Y_a is sensible conceptually, estimating $E(Y_a)$ non-parametrically is hard for continuous exposures

- To define and estimate causal effects in the presence of **violations of the positivity assumption**:
 - IPSIs satisfy positivity by design
 - MTPs can also be arranged to satisfy positivity by definition (but require some knowledge about the support of A)

Identification using reweighting

The idea of inverse probability weighting also applies to estimation of these parameters.

Each individual needs to be reweighted to account for the number of individuals it would represent in a hypothetical population where the intervention has been performed

Specifically, define the density ratio

$$r(A, W) = \frac{g_{A,\delta}(A | W)}{g_A(A | W)}$$

The **reweighting identification formula** is given by

$$E[Y(A_\delta)] = E[r(A, W) \times Y]$$

Simple estimation procedures

The reweighted identification formula suggests a simple estimation strategy.

First, we construct an estimate of the density ration $r(A, W)$.

Then, we can compute an **IPTW estimator** of $E[Y(A_\delta)]$ as

$$\psi_{n,IPTW} := \frac{1}{n} \sum_{i=1}^n r_n(A_i, W_i) Y_i .$$

There are at least two possible strategies to estimate the density ratio

$$r(A, W) = \frac{g_{A,\delta}(A | W)}{g_A(A | W)}$$

- 1 Construct estimates of the densities $g_A(A, W)$ and $g_{A,\delta}(A, W)$, and plug them into the above definition of $r(A, W)$.
 - This may not be easy to do with data-adaptive estimators: the machine and statistical learning literatures have only few methods for conditional density estimation
- 2 **Estimate the density ratio directly by recasting the problem as a classification problem** (Details in Díaz et al. 2020 and in the lab for this chapter)

The G-computation formula suggests another natural estimation strategy

$$E[Y(A_\delta)] = E[\bar{Q}_Y(A_\delta, W)],$$

- 1 **STEP 1:** Regress Y on A and $W \rightarrow \bar{Q}_{Y,n}$
- 2 **STEP 2:** Predict under the intervention $\rightarrow \bar{Q}_{Y,n}(A_\delta, W)$
- 3 **STEP 3:** Average the predictions across the sample

$$\psi_{n,gcomp} := \frac{1}{n} \sum_{i=1}^n \bar{Q}_{Y,n}(A_{\delta,i}, W_i)$$

As before, a hybrid between the G-computation and IPTW estimator can be constructed, and it enjoys improved properties

The **one-step estimator** (akin to AIPTW) is given by

$$\psi_{n,os} := \psi_{n,gcomp} + \frac{1}{n} \sum_{i=1}^n r_n(A_i, W_i) [Y - \bar{Q}_{Y,n}(A_i, W_i)]$$

One-step estimator in R

The one-step estimator enjoys the following improved properties:

- It hits the right target if **either** r **or** \bar{Q}_Y is estimated well
- Valid confidence intervals (even when flexible regression is used) can be constructed if **both** r **and** \bar{Q}_Y are estimated well

Improved estimation procedures

A targeted minimum loss based estimator can be constructed in the following additional steps

- 1 Fit a tilting logistic regression model

$$\text{logit } Q_Y(A, W) = \text{logit } Q_{Y,n}(A, W) + \varepsilon,$$

by:

- Regressing Y with an intercept-only logistic regression model,
 - where the variable $\text{logit } Q_{Y,n}(a, w)$ is taken as an offset, and
 - The regression is fit using weights $r_n(A_i, W_i)$
- 2 Computed the updated outcome predictions under the intervention as the G-computation estimator

$$\tilde{Q}_{Y,n}(A_\delta, W) = \text{expit} \{ \text{logit } \bar{Q}_{Y,n}(A_\delta, W) + \varepsilon_n \}$$

- 3 Compute the TMLE as

$$\psi_{n,tMLE} := \frac{1}{n} \sum_{i=1}^n \tilde{Q}_{Y,n}(A_{\delta,i}, W_i)$$

These estimators are implemented in several R packages:

- `lmtp` (on CRAN)
 - Implements cross-fitting for improved properties
 - Longitudinal data
- `txshift`
- `tmle3`

Examples in the lab for this chapter

A simulation study illustrating the properties of the TMLE

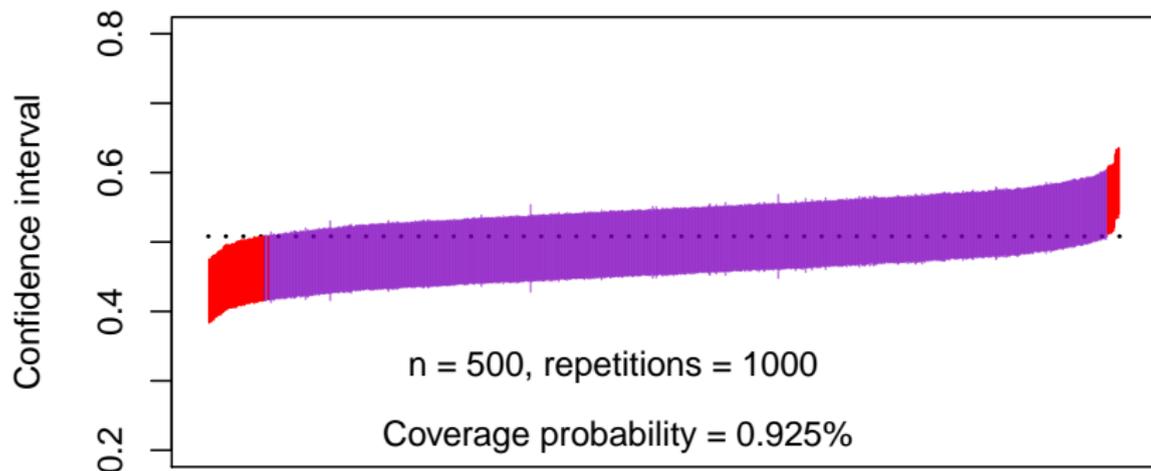
We simulate data as follows

$$\begin{aligned} Y | A = a, W = w &\sim \text{Bernoulli}(\text{expit}(1 - a + w^2)) \\ A | W = w &\sim \text{Normal}(1 - 0.5 \times \log(|w|), 1) \\ W &\sim \text{Normal}(0, 1) \end{aligned}$$

- Sample size $n = 500$
- Simulate $m = 1000$ datasets
- Run `lmtp` using a Super Learner. The library contains:
 - MARS and logistic regression, or
 - only logistic regression
- Look at the bias and coverage of confidence intervals across the $m = 1000$ datasets

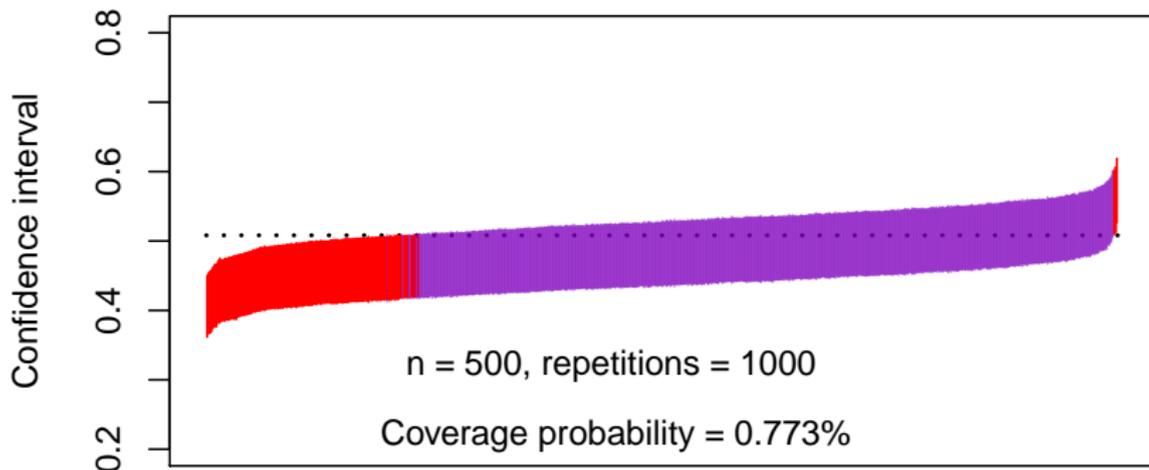
A simulation study illustrating the properties of the TMLE

`c('SL.glm', 'SL.earth')`



A simulation study illustrating the properties of the TMLE

SL.glm



References:

Kennedy, Edward H. "Nonparametric causal effects based on incremental propensity score interventions." *Journal of the American Statistical Association* 114.526 (2019): 645-656.

Young, Jessica G., Miguel A. Hernán, and James M. Robins. "Identification, estimation and approximation of risk under interventions that depend on the natural value of treatment using observational data." *Epidemiologic methods* 3.1 (2014): 1-19.

Díaz, Iván, and Mark van Der Laan. "Population intervention causal effects based on stochastic interventions." *Biometrics* 68.2 (2012): 541-549.

Díaz, Iván, et al. "Non-parametric causal effects based on longitudinal modified treatment policies." *arXiv preprint arXiv:2006.01366* (2020).

An effect decomposition for stochastic interventions

Assume we want to decompose the total effect

$$E[Y(A_\delta)] - E[Y]$$

into:

- A direct effect operating through $A \rightarrow Y$
- An indirect effect operating through $A \rightarrow M \rightarrow Y$

An effect decomposition for stochastic interventions

As before, we have the following definitions:

- $Y(a)$ is the counterfactual obtained by setting $A = a$
- $Y(a, m)$ is the counterfactual obtained by setting $A = a$ and $M = m$

Notice that an intervention setting A equal to A_δ induces two counterfactual variables:

- $M(A_\delta)$
- $Y(A_\delta) = Y(A_\delta, M(A_\delta))$

An effect decomposition for MTPs

$$\begin{aligned}\psi(\delta) &= E[Y(A_\delta) - Y] \\ &= \underbrace{E[Y(A_\delta, M(A_\delta)) - Y(A_\delta, M)]}_{\text{indirect effect}} + \underbrace{E[Y(A_\delta, M) - Y(A, M)]}_{\text{direct effect}}.\end{aligned}$$

- We have discussed estimation of $E[Y(A_\delta)] = E[Y(A_\delta, M_\delta)]$
- As before, estimation of $E(Y) = E[Y(A, M)]$ is straightforward
- In what follows we focus on **identification and estimation of $E[Y(A_\delta, M)]$**

An effect decomposition for MTPs

- Identification and estimation methods developed in Díaz and Hejazi, 2020
- The `medshift` R package (see lab for this chapter) provides an implementation of the relevant estimators

Identification of the counterfactual mean $E[Y(A_\delta, M)]$

We need a modified version of the usual two assumptions:

- **Positivity:**

- If $g_{A,\delta}(a | w) > 0$ then $g_A(a | w) > 0$
- If $g_M(m | w) > 0$ then $g_M(m | a, w) > 0$

- **Randomization:**

- $Y(a, m) \perp A | W$
- $Y(a, m) \perp M | (A, W)$

Under these assumptions we have

$$E[Y(A_\delta, M)] = E[\bar{Q}_Y(M, d(A, W), W)]$$

Identification of the counterfactual mean $E[Y(A_\delta, M)]$

Identification if the intervention is defined as $A_\delta = d(A, W)$

$$\begin{aligned} E[Y(A_\delta, M) \mid M = m, A = a, W = w] &= E[Y(d(a, w), m) \mid M = m, A = a, W = w] \\ &= E[Y(d(a, w), m) \mid M = m, A = d(a, w), W = w] \\ &= E[Y \mid M = m, A = d(a, w), W = w] \\ &= \bar{Q}_Y(m, d(a, w), w), \end{aligned}$$

where we have defined

$$E(Y \mid M = m, A = a, W = w) := \bar{Q}_Y(m, a, w)$$

Averaging with respect to the distribution of (M, A, W) yields

$$E[Y(A_\delta, M)] = E[\bar{Q}_Y(M, d(A, W), W)]$$

Identification of the counterfactual mean $E[Y(A_\delta, M)]$ for MTPs

- Let $\tilde{W} = (W, M)$
- The identification formula becomes

$$E \left[\bar{Q}_Y(d(A, \tilde{W}), \tilde{W}) \right],$$

where

$$\bar{Q}_Y(a, \tilde{w}) = E[Y \mid A = a, \tilde{W} = \tilde{w}]$$

- This is identical to the formula that we studied in the first part of this chapter
- Thus, for purposes of estimation, we can forget about the nature of M as a mediator and W as a confounder and proceed using the methods already discussed.

Identification of mediational effects

The **indirect effect** is thus identified as

$$E [\bar{Q}_Y(d(A, W), W) - \bar{Q}_Y(M, d(A, W), W)]$$

The **direct effect** is identified as

$$E [\bar{Q}_Y(M, d(A, W), W) - \bar{Q}_Y(M, A, W)]$$

A note on identification assumptions

Let us revisit the **randomization assumption**:

- $Y(a, m) \perp A \mid W$
- $Y(a, m) \perp M \mid (A, W)$

As before, this assumption precludes intermediates confounders Z

However, unlike the cross-world assumption necessary for identification of the NIE/NDE, **this assumption can be satisfied by design** if the study randomizes both the mediator and the treatment.

Stochastic effect decomposition in the presence of intermediate confounders

- Assume a confounder Z of the relation $M \rightarrow Y$ is affected by A
- The above effects are unidentified in this case
- A solution is to use a randomized mediator, as before
- Let G denote a random draw distributed as $M \mid (A, W)$
- Let G_δ denote a random draw distributed as $M(A_\delta) \mid (A_\delta, W)$

Then we can use the following effect decomposition:

$$E[Y(A_\delta, G_\delta) - Y(A, G)] = \underbrace{E[Y(A_\delta, G_\delta) - Y(A_\delta, G)]}_{\text{indirect effect}} + \underbrace{E[Y(A_\delta, G) - Y(A, G)]}_{\text{direct effect}}$$

References:

Hejazi, Nima S., et al. "Nonparametric causal mediation analysis for stochastic interventional (in) direct effects." arXiv preprint arXiv:2009.06203 (2020).

Díaz, Iván, and Nima S. Hejazi. "Causal mediation analysis for stochastic interventions." *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 82.3 (2020): 661-683.