
Causal inference with multiple concurrent medications: a comparison of methods and an application in multidrug-resistant tuberculosis

Arman Alam Siddique¹, Mireille E. Schnitzer^{*,2}, Asma Bahamyirou², Guanbo Wang³, Timothy H. Holtz⁴, Giovanni Battista Migliori⁵, Giovanni Sotgiu⁶, Neel R. Gandhi⁷, Mario H. Vargas^{8,9}, Dick Menzies^{10,11}, and Andrea Benedetti^{3,10,11}

Abstract

This paper investigates different approaches for causal estimation under multiple concurrent medications. Our parameter of interest is the marginal mean counterfactual outcome under different combinations of medications. We explore parametric and non-parametric methods to estimate the generalized propensity score. We then apply three causal estimation approaches (inverse probability of treatment weighting, propensity score adjustment, and targeted maximum likelihood estimation) to estimate the causal parameter of interest. Focusing on the estimation of the expected outcome under the most prevalent regimens, we compare the results obtained using these methods in a simulation study with four potentially concurrent medications. We perform a second simulation study in which some combinations of medications may occur rarely or not occur at all in the dataset. Finally, we apply the methods explored to contrast the probability of patient treatment success for the most prevalent regimens of antimicrobial agents for patients with multidrug-resistant pulmonary tuberculosis.

Keywords

Causal Inference, Concurrent Medications, Generalized Propensity Score, Machine Learning, Multidrug-Resistant Tuberculosis, Targeted Maximum Likelihood Estimation.

1 Introduction

Polypharmacy is the intake of multiple medications, potentially more than medically necessary, at the same time. Apart from the increased costs for multiple medications, the degradation of quality of life, the possibility of interactions between those medications, and adverse drug reactions¹, make polypharmacy an important area of research.

The concurrent usage of multiple medications is necessary for some diseases. Multidrug-resistant tuberculosis (MDR-TB), with almost 500 000 new cases in 2016² and a 45% mortality rate worldwide³, is defined as a disease caused by strains of *Mycobacterium tuberculosis* that are resistant to at least the two most effective drugs, isoniazid and rifampin, used to treat tuberculosis. Patients with MDR-TB are treated with multiple alternative antimicrobial agents in order to cure the infection and prevent further drug resistance (or to prevent the selection of drug resistant strains of *M. tuberculosis*). Current guidelines recommend the simultaneous usage of five or more antimicrobial agents depending on the therapeutic phase and drug resistance pattern⁴. A systematic review published in 2012 identified international studies that investigated associations between

¹ Department of Statistics, McMaster University, Hamilton, Canada

² Faculty of Pharmacy, Université de Montréal, Montreal, Canada

³ Department of Epidemiology, Biostatistics & Occupational Health, McGill University, Montreal, Canada

⁴ Division of Global HIV and TB, Centers for Disease Control and Prevention, New Delhi, India.

⁵ World Health Organization Collaborating Centre for Tuberculosis and Lung Diseases, Fondazione S. Maugeri, Tradate, Italy

⁶ Clinical Epidemiology and Medical Statistics Unit, Department of Medical, Surgical and Experimental Sciences, University of Sassari, Sassari, Italy

⁷ Rollins School of Public Health and Emory School of Medicine, Emory University, Atlanta, USA

⁸ Departamento de Investigación en Hiperreactividad Bronquial, Instituto Nacional de Enfermedades Respiratorias, Mexico City, Mexico

⁹ Unidad de Investigación Médica en Enfermedades Respiratorias, Instituto Mexicano del Seguro Social, Mexico City, Mexico

¹⁰ Respiratory Epidemiology and Clinical Research Institute, McGill University Health Centre, Montreal, Canada

¹¹ Department of Medicine, McGill University, Montreal, Canada

Corresponding author:

Mireille E. Schnitzer, Faculty of Pharmacy, Université de Montréal, Montreal, Québec, Canada, H3C3J7

Email: mireille.schnitzer@umontreal.ca

different treatments and treatment outcomes of MDR-TB⁵. The combination of individual patient data from these studies is currently the greatest resource for evaluating medication effectiveness in MDR-TB. However, with patients taking as many as 7 antimicrobial agents concurrently⁵, and the data containing 15 different antimicrobial agents overall, the analysis presents a challenge for the application of causal inference methods.

Many causal estimation techniques for binary treatments use the propensity score, defined as the probability of receiving one of the two treatment options. In the case where multiple treatments are available, Imbens⁶ extended this framework by defining the *generalized propensity score* (GPS) as the probability of receiving a specific treatment. Imbens⁶, Imai and Van Dyk⁷, and Lopez and Gutman⁸ developed various techniques reliant on the GPS for the estimation of causal effects. Further, McCaffrey *and others*⁹ proposed using generalized boosted models for the estimation of the GPS for multiple treatments.

In this paper, we explore methods to estimate the relative effects of taking multiple medications. The previous methods cited above primarily estimated the effects of continuous (such as medication dose) or low-dimensional categorical treatment options. In contrast, we are interested in the setting where patients may take more than one medication of interest concurrently, resulting in a potentially large number of possible drug combinations, many of which may not be observed in the data.

In order to approach this problem, we take the exposure to be a categorical variable of *regimens*, where regimen refers to a specific combination of medications (perhaps taken over a pre-specified period). We then employ various machine learning algorithms for the estimation of the GPS. We provide short introductions for these machine learning algorithms along with several causal estimation procedures in Section 2. We present a simulation study in Section 3 in order to compare the appropriateness of each method. In Section 4, we present an application of these methods for the MDR-TB data in which we provide estimates of the expected rates of treatment success (with outcome defined by the World Health Organization⁴) for the 10 most prevalent regimens in the combined dataset of Ahuja *and others*⁵.

2 Methods

In order to estimate the causal effects of multiple medications, we propose to estimate the GPS, defined as the probability of taking a specific regimen conditional on covariates. To this end, we investigate the usage of different machine learning algorithms for the GPS. Further, in order to estimate the causal contrasts, we employ Inverse Probability of Treatment Weighting¹⁰, Propensity Score Adjustment¹¹ and Targeted Maximum Likelihood Estimation^{12,13}, all of which use the GPS. We also investigate G-Computation¹⁴, which exclusively uses a model for the outcome conditional on medications taken and covariates in order to estimate an effect of interest.

2.1 General Notation

The observed data O_i include a vector of covariates, $X_i = \{X_{ij}; j = 1, \dots, J\}$, and a univariate outcome, Y_i where $i = 1, \dots, n$ indexes the set of subjects. We consider a fixed set of K potential medications that all patients in the study are hypothetically eligible for. For any patient i , the binary variable A_i^k indicates exposure to medication $k \in \{1, \dots, K\}$. We define $C_i = (A_i^1, \dots, A_i^K)$ as the set of treatments being taken by patient i . We denote R_i as a categorical variable corresponding to the observed regimen for patient i , represented by the combination of treatments C_i . For each individual, R_i corresponds to one of the 2^K different possible regimens. We denote a specific fixed regimen as r and the corresponding vector of binary elements as c^r . We also define B_i^r as an indicator for the regimen r , i.e. if patient i took regimen r , then $B_i^r = 1$. Clearly, C_i , R_i , and B_i^r contain the same information, but we require these definitions in order to describe the proposed models. We drop the i subscript when referring to a random draw of a variable from the population.

The goal of the analysis is to estimate $\mathbb{E}(Y^r)$, which is also equivalent to $\mathbb{E}(Y^{c^r})$, where Y_i^r or $Y_i^{c^r}$ represents the potential outcome of subject i had they received an intervention corresponding with a treatment regimen r . We may then contrast different regimens by comparing their respective estimated values of $\mathbb{E}(Y^r)$. In MDR-TB example, the binary outcome is defined as treatment success (the treatment was completed and cured the infection) versus failure (patient still tested culture positive for MDR-TB, died, or defaulted on treatment/were lost to follow-up). The goal of the application was therefore taken to be the estimation

of the probability of treatment success under a given regimen of antimicrobial agents. The regimens with the higher probabilities of treatment success may then be interpreted as having greater effectiveness than those with a lower probability.

2.2 Estimation of the Generalized Propensity Score

The propensity score¹⁵ is defined as the probability of receiving a treatment conditional on covariates. When dealing with a binary treatment where $C \in \{0,1\}$, the propensity score can be mathematically expressed as

$$g(X) = Pr(C = 1|X).$$

With multiple treatments, the propensity score was extended to the GPS⁶ defined as

$$g(r, X) = Pr(R = r|X) = Pr(C = c^r|X),$$

the probability of receiving a given regimen r . We use multi-class classification, with classes corresponding to regimens, in order to estimate the GPS. Multi-class classification is the fitting of models for different classes in the dataset where the classes are mutually exclusive. In this subsection we provide basic descriptions of support vector machines, softmax regression (i.e. multinomial regression), and generalized boosted models, which we later use to estimate the GPS.

2.2.1 Support Vector Machines Support Vector Machines (SVMs) (Hastie and others¹⁶, Chapter 12), a supervised learning approach, have been proposed as a method for multi-class classification and have been identified as one the most important research topics in the field of machine learning¹⁷. Computationally efficient, SVMs use hyperplanes to delineate a particular class by identifying the most influential observations in the determination of the boundaries between the classes. These observations are also known as the support vectors. The main aim of SVMs is to find a maximum margin hyperplane, where margin corresponds to the distance between the hyperplane and closest elements on either side of the hyperplane.

For the pairwise classification of two different regimens, say r_1 and r_2 , Soft-Margin SVMs¹⁸ construct a hyperplane $\{X; f(X) = w^T X + b = 0\}$, with the constraint $\{I(R_i = r_1) - I(R_i = r_2)\}(w^T X_i + b) \geq 1 - \zeta_i$, for all $i = 1, \dots, n$ where the $\zeta_i \geq 0$ are called “slack variables” and $I(\cdot)$ is the indicator function. If

$\zeta_i = 0$ for all $i = 1, \dots, n$, this would imply that the hyperplane would be able to perfectly separate and classify the data. The slack variables therefore allow for misclassification.

The parameters w , b and ζ_i are estimated by minimizing a loss function $F(w, b, \zeta)$ over w and b subject to the above constraints. This loss function is given by:

$$F(w, b, \zeta) = \frac{\|w\|^2}{2} + C \sum_{i=1}^n \zeta_i,$$

where C is a constant which maintains the trade-off between the training error and the margins (a smaller C allows for a smoother boundary $f(X)$). The function $F(w, b, \zeta)$ is minimized using optimization methods with Lagrangian multipliers¹⁶.

We apply the default settings of the function `svm` in the `e1071` R package¹⁹ for the implementation of SVMs. In particular, this function uses One-Vs-One classification²⁰ (i.e. constructs boundaries for each pair of classes separately, and the final classification for each observation is determined by which class is most frequently selected), sets $C = 1$, and applies a non-linear basis expansion with a radial kernel (Hastie *and others*¹⁶, Section 12.3). Finally, the probability of class membership (following a given regimen r) is estimated by fitting a logistic regression of $R = r$ on the boundary values $f(X)$ computed for each pairwise comparison^{21,22}.

2.2.2 Softmax Regression Softmax regression²³, a common classification method, is equivalent to multinomial logistic regression. We restrict the probability for a patient to be treated with regimen r as

$$Pr(R_i = r | X_i, \Phi) = \frac{\exp(\phi_r^T X_i)}{\sum_{l=1}^{2^K} \exp(\phi_l^T X_i)}$$

The model parameters $\phi_r \in \mathbb{R}^{J+1}$, $r \in \{1, \dots, 2^K\}$, with J corresponding to the number of covariates present in the model, are stacked together to form Φ , a matrix of dimension $2^K \times (J + 1)$ with entries $\Phi_{r,j}$. The parameter matrix Φ is then estimated by minimizing the loss function $L(\Phi)$ (corresponding with the

negative quasi log-likelihood), which is given by:

$$L(\Phi) = \sum_{i=1}^n \sum_{r=1}^{2^K} I(R_i = r) \log \frac{\exp(\phi_r^T X_i)}{\sum_{l=1}^{2^K} \exp(\phi_l^T X_i)}$$

For implementation, we use the `softmaxreg` package²⁴ in R.

2.2.3 Generalized Boosting Generalized Boosted Models (GBMs) (Hastie and others¹⁶, Chapter 10) are machine learning algorithms that build up an additive model using multiple classification trees. Classification trees (Hastie and others¹⁶, Chapter 9) create a piecewise model for a treatment by learning which sequential splits in the covariates most improve prediction of the treatment. Boosting generates a sequence of trees while upweighting the observations that were misclassified by the previous trees. Finally, the predictions from the individual trees are combined using an error-weighted majority vote.

Implementations of GBMs have been proposed to estimate the GPS for multiple treatments. To prevent overfitting, one needs to identify the total number of trees to use. McCaffrey and others⁹ propose to select the number of trees by comparing the values of the covariates in the GPS-weighted treatment group versus the entire sample. A good “balance” means that covariate distributions are similar between these groups. The number of trees can be chosen by satisfying a criterion such as the Absolute Standard Bias (ASB), which compares the standardized difference in covariate means between groups, or the Kolmogorov-Smirnov (KS) Statistic, which compares the empirical distributions. In addition to the number of trees, the tuning parameters include a shrinkage term (learning rate) for the GBM, the minimum number of observations in the trees’ terminal nodes, and the depth of interactions (indicating the maximum number of splits the algorithm performs on a tree after the initial split) included in the model, all of which are important in order to properly smooth the model. We estimate the GPS for each regimen separately using the `twang` package²⁵ in R.

2.3 Causal Estimation Methods

After obtaining the GPS, we aim to estimate $E(Y^r)$, where Y^r is the potential outcome of an arbitrary patient under regimen r . In order to obtain an estimate of $E(Y^r)$, one may choose from various causal estimation methods, several of which we describe in this subsection. Causal estimation methods adjust for

the confounders (roughly, those pre-treatment variables X that are related to both treatment regimen and Y) in order to produce estimates of the marginal parameter $\mathbb{E}(Y^r)$. These causal estimation methods rely on several assumptions⁶, including 1) positivity: the probability of receiving any regimen r conditional on the confounders, X , should be a non-zero quantity for all subjects; 2) consistency: for any patient i taking regimen $R_i = r$, the counterfactual outcome for patient i under r is the observed outcome of the patient; and 3) conditional exchangeability: the observed covariates should be sufficient to satisfy conditional independence between the regimens and the potential outcomes. Since we have 2^K different regimens, some of which may not at all be observed in the data, the assumption of positivity is very likely to fail (either empirically or theoretically) for some regimens. This would imply that without additional extrapolation, we would not be able to estimate $\mathbb{E}(Y^r)$ for those regimens. In the following, we only estimate the parameter of interest for prevalent regimens.

2.3.1 G-Computation G-computation is a causal estimation method proposed by Robins¹⁴ that can be used for the estimation of $\mathbb{E}(Y^r)$. The algorithm for G-Computation²⁶ is as follows:

Algorithm 1 G-Computation for $E(Y^r)$

- 1: Fit an outcome model for $\mathbb{E}(Y|R, X)$ using the available data, defined as $Q(R, X)$. We then compute predictions of the conditional expectations under the regimen r for every subject. In our context, one may use a model $Q^{(a)}(R, X)$ that is conditional on the regimens directly (i.e. subsetting on $B^r = 1$ or taking the indicators B^r as covariates) or an alternative $Q^{(b)}(C, X)$ that is conditional on the medications (taking the A^k as covariates).
- 2: For each observation, predict the value of $Q_n(r, X_i) = \mathbb{E}_n(Y|R = r, X_i)$ using the above obtained model where a subscript n denotes an estimate of the quantity.
- 3: The G-computation estimate of $\mathbb{E}(Y^r)$ is thus given by:

$$\psi_{n,G-comp}^r = \frac{1}{n} \sum_{i=1}^n Q_n(r, X_i).$$

The unbiasedness of G-computation relies on the correct specification of the outcome model.

2.3.2 Inverse Probability of Treatment Weighting Inverse Probability of Treatment Weighting (IPTW)¹⁰ is an approach for the estimation of $\mathbb{E}(Y^r)$ using the propensity score. The algorithm for performing IPTW is as follows:

Algorithm 2 Inverse Probability of Treatment Weighting for $E(Y^r)$

- 1: Estimate the GPS for each regimen, $g_n(r, X_i) = Pr_n(R = r|X_i)$.
 - 2: Obtain the weight $w_n(r, X_i) = I(R_i = r)/g_n(r, X_i)$ for each observation, which is only non-zero for subjects who took the regimen of interest, r .
 - 3: Run a linear regression model of Y on an intercept, with weights $w_n(r, X)$.
-

The resulting estimate of the intercept is our IPTW estimate, $\psi_{n,IPTW}^r$. The consistency of IPTW relies on the correct specification of the propensity score model. In order to calculate the variance of the resulting estimate we use the `sandwich` package²⁷ in R, which is used for calculating robust variance estimates (that take into account the uncertainty in the propensity score). One could alternatively use the non-parametric bootstrap to estimate the variance, but this may be excessively time-consuming when the GPS is estimated with a machine learning method.

2.3.3 Propensity Score Adjustment Propensity Score Adjustment (PSA) is a causal estimation method that relies on the specification of the propensity score model in addition to a model for the outcome, conditional on the propensity score and treatment. The propensity score is a balancing statistic, that is, given the propensity score, the potential outcome is conditionally independent of the treatment¹¹. For a single binary treatment $C \in \{0, 1\}$, one might use the following model:

$$\mathbb{E}(Y|C, g(X)) = \theta_0 + \theta_1 C + \theta_2 g(X),$$

where θ_1 can also be written as $\theta_1 = \mathbb{E}(Y|C = 1, g(X)) - \mathbb{E}(Y|C = 0, g(X))$. The estimate $\hat{\theta}_1$ is obtained using least squares and is an unbiased estimate of $\mathbb{E}(Y|C = 1, g(X)) - \mathbb{E}(Y|C = 0, g(X)) = \mathbb{E}(Y^1 - Y^0)$ if the propensity score and the outcome regression model are correctly specified and if the causal assumptions hold. However, if the expected outcome is not linearly dependent on the propensity score or if the propensity score model is incorrectly specified, then the ordinary

least squares estimate of θ_1 is a biased estimator of $\mathbb{E}(Y^1 - Y^0)$ ¹¹. If a non-linear model is used, the above result may not be applicable, since θ_1 for this case might correspond with a conditional parameter (and estimation would therefore be biased for the marginal contrast between the potential outcomes).

This method of estimation can also be extended to the case with multiple treatments⁶. For our setting, we propose the following algorithm.

Algorithm 3 Propensity Score Adjustment for $E(Y^r)$

- 1: Fit a model $Q^{(1)}(R, g(r, X))$ (conditional on the regimen indicators, B^r) or $Q^{(2)}(C, X)$ (conditional on the treatment indicators, A^k) for $\mathbb{E}(Y|R, g(r, X))$.
- 2: Using the model fit, obtain predictions:

$$Q_n^{(1)}(r, g_n(r, X)) = \mathbb{E}_n(Y|B^r = 1, g_n(r, X)) \text{ or}$$

$$Q_n^{(2)}(c^r, g_n(r, X)) = \mathbb{E}_n(Y|C = c^r, g_n(r, X))$$

- 3: The estimates of $\mathbb{E}(Y^r)$ are then given as

$$\psi_{n,PSA(I)}^r = \frac{1}{n} \sum_{i=1}^n Q_n^{(1)}(r, g_n(r, X_i)), \text{ and}$$

$$\psi_{n,PSA(II)}^r = \frac{1}{n} \sum_{i=1}^n Q_n^{(2)}(c^r, g_n(r, X_i)).$$

2.3.4 Targeted Maximum Likelihood Estimation Targeted Maximum Likelihood Estimation¹³ (TMLE) is a semi-parametric estimation technique that produces doubly robust and locally efficient plug-in estimators. In our situation, TMLE invokes a two-step process that first produces estimates of the conditional expectation of the outcome under a fixed regimen (as in the first step in G-Computation) and then updates these initial estimates²⁸. The update procedure uses the propensity score and is designed to reduce the bias in the estimate of the causal parameter of interest. An algorithm for the computation of TMLE for the multiple medication case with target parameter $\mathbb{E}(Y^r)$ is described below.

Algorithm 4 Targeted Minimum Loss-Based Estimation for $E(Y^r)$

- 1: First, fit an outcome model and generate estimates of the conditional expectation under the fixed regimen r , denoted $Q_n(r, X)$. We may use $Q_n^{(a)}(r, X)$ or $Q_n^{(b)}(c^r, X)$ as described in Section 2.3.1.
- 2: Define weights $w_n(r, X) = I(R = r)/g_n(r, X)$.
- 3: Regress Y on 1 with offset $\text{logit}\{Q_n(r, X)\}$ and weights $w_n(r, X)$. Denote the estimate of the intercept term by $\hat{\epsilon}$.
- 4: Compute the updated estimate, $Q_n^*(r, X)$, which is given by:

$$\text{logit}(Q_n^*(r, X)) = \text{logit}(Q_n(r, X)) + \hat{\epsilon}.$$

- 5: The TMLE estimate for $\mathbb{E}(Y^r)$ is then given by:

$$\psi_{n, TMLE}^r = \frac{1}{n} \sum_{i=1}^n Q_n^*(r, X_i).$$

The double robustness property of this TMLE means that, unlike the propensity score adjustment method, the TMLE is a consistent estimator if either $\mathbb{E}(Y | R = r, X)$ or $g(r, X)$ is consistently estimated. For the approximation of the estimation standard error, one may use the *efficient influence function* (EIF)²⁹, corresponding to the first order expansion of the estimator:

$$EIF^r(Q, g)(O) = (Y - Q(r, X)) \frac{I(R = r)}{g(r, X)} + Q(r, X) - \psi_{TMLE}^r.$$

In large samples, the variance of the estimator will correspond to the sample variance of the estimated EIF. Therefore, the 95% confidence interval for $\psi_{n, TMLE}^r$ can be estimated by $\psi_{n, TMLE}^r \pm 1.96 \sqrt{(\sigma_{n, TMLE}^r)^2/n}$, where $(\sigma_{n, TMLE}^r)^2$ denotes the sample variance of $EIF^r(Q_n, g_n)(O_i)$.

3 Simulation Study

In order to evaluate the appropriateness of the above causal estimators paired with each GPS method, we contrast their performance in a Monte Carlo simulation study. We first describe the data generating mechanisms. We estimate the expected counterfactual outcomes under the most prominent regimens. We

compare the performance of several implementations of G-computation and then of each causal estimator that uses the GPS. In the Supplementary Materials, we perform a second simulation study with a larger number of medications, often leading to more regimens than subjects in the sample. For this second scenario, we evaluate a data subsetting method that can greatly reduce computational time.

3.1 Data Generation

Full details of the data generation are given in Section 1 of the Supplementary Materials.

We independently generate 12 baseline variables $X_{ij}, j = 1, \dots, 12$ from a standard uniform distribution i.e. $X_{ij} \sim U(0, 1)$. We also generate 4 dichotomous treatment indicators, $A_i^k, k = 1, 2, 3, 4$, conditional on a subset of the baseline variables. In addition, A^1 and A^2 are generated as positively correlated as are A^3 and A^4 , and all other treatments pairs are independent. Specifically, a patient is more likely to take medication 1 if they are also taking medication 2 (and vice versa), and similarly for medications 3 and 4. A binary outcome Y_i is generated using a logistic model conditional on the X_{ij} s and A_i^k s with first-order interactions (including treatment-treatment, covariate-covariate, and covariate-treatment interactions). As subjects can take up to 4 medications, there are $2^4 = 16$ possible regimens. The two most likely regimens (on average) are regimen 1 (1,1,0,0) and regimen 2 (1,1,1,1) and are defined as the regimens of interest. The true propensity score $Pr(A^1 = a_1, A^2 = a_2, A^3 = a_3, A^4 = a_4 | X)$ in this case can be factorized as:

$$Pr(A^1 = a_1, A^2 = a_2, A^3 = a_3, A^4 = a_4 | X) = Pr(A^1 = a_1, A^2 = a_2 | X) \cdot Pr(A^3 = a_3, A^4 = a_4 | X)$$

The true values of $\mathbb{E}(Y^r)$ are 0.61 and 0.57 for regimens 1 and 2 respectively.

3.2 Comparison of Outcome Regression Models

Since propensity score adjustment and TMLE both use a model for the outcome, we first evaluate the performance of six implementations of G-Computation to see whether each outcome model produces biased effects of $E(Y^r)$. We fit the following outcome models with logistic regressions: 1) by regimen, subsetting on $B^r = 1$ for each r of interest, and 2) by treatment, adjusting for the treatment indicators

A^k , $k = 1, 2, 3, 4$ in the regression. For the latter case, we fit the outcome models without interactions (taking the main terms of A^k only) and then with first-order interactions between the A^k . We apply these three approaches to G-Computation both with and without adjustment for the baseline covariates as main terms.

We generated 1 000 datasets of sample size $n = 500$ and 1 000, respectively. Table 1 gives the mean estimates and Monte Carlo standard errors for each implementation. For regimen 1, the G-computation estimate had little bias when adjusting by regimen or by treatment with first order interactions, regardless of the adjustment for X_{ij} as main terms. However, it was substantially biased when fit with treatment main terms only, regardless of adjustment for X_{ij} . For regimen 2, the G-computation estimate was unbiased when adjusting by regimen or by treatment with first order interactions but only when also adjusting for confounding by X_{ij} . It was biased when not adjusting for confounding and when the treatment interactions were not included. The standard error was lower for the larger sample size but the bias remained steady.

Table 1. Monte Carlo mean estimates and standard errors for different implementations of G-Computation. The true value for regimen 1 is $\mathbb{E}(Y^1) = 0.61$ and the true value for regimen 2 is $\mathbb{E}(Y^2) = 0.57$. Q_n corr indicates whether the outcome model includes the true treatment-treatment interactions.

	Q_n corr	$n = 500$		$n=1\ 000$	
		Reg 1	Reg 2	Reg 1	Reg 2
Unadjusted					
By Regimen	Y	0.63(0.05)	0.62(0.07)	0.63(0.03)	0.62(0.04)
By treatment (main terms)	N	0.48(0.05)	0.83(0.03)	0.48(0.03)	0.83(0.03)
By treatment (first order interactions)	Y	0.63(0.05)	0.62(0.07)	0.63(0.03)	0.62(0.04)
Adjusted for X_{ij}					
By Regimen	Y	0.64(0.04)	0.57(0.07)	0.64(0.03)	0.57(0.05)
By treatment (main terms)	N	0.47(0.04)	0.81(0.03)	0.47(0.03)	0.82(0.03)
By treatment (first order interactions)	Y	0.62(0.05)	0.58(0.06)	0.62(0.03)	0.58(0.04)

3.3 Comparison of methods

The implementations of causal estimators that are evaluated in this section are:

- IPTW, using a weighted linear regression model (Section 2.3.2);

- PSA(I), propensity score adjustment with a logistic regression to estimate $Q^{(1)}$ conditional on regimen (Section 2.3.3);
- PSA(II), propensity score adjustment with a logistic regression to estimate $Q^{(2)}$ conditional on treatments as main terms (Section 2.3.3);
- TMLE(I), using a logistic regression to model the outcome conditional on regimen and baseline covariates, i.e. $Q^{(a)}$ (Section 2.3.4);
- TMLE(II), using a logistic regression to model the outcome conditional on treatments and baseline covariates, i.e. $Q^{(b)}$ (Section 2.3.4).

The GPS for each regimen of interest was estimated using the three approaches in Section 2.2. When fitting GBMs for each regimen R_i , we chose values of the tuning parameters that optimized the balance between the pre-treatment covariates in R_i and the pooled sample of all the other regimens for 5 simulated datasets using the `plots` function in `twang`. The maximum number of iterations in the Softmax regression was set to 100 with the default learning rate of 0.05 and the tuning parameters for SVMs were similarly assigned the default values.

We drew 1 000 samples of sizes $n = 500$ and 1 000, respectively. Table 2 gives the mean estimates and Monte Carlo standard errors for the top 2 occurring regimens in our simulated data. The numbers of subjects exposed to each of these regimens varied by sample and are given in Section 3 of the Supplementary Materials. TMLE performed well when implemented with SVMs, Softmax regression, and GBMs. IPTW and PSA(I) performed well with Softmax regression but were biased with SVMs and GBMs for the second regimen, likely due to the suboptimal convergence rate of these nonparametric GPS methods³⁰. The estimates of PSA(I) and IPTW with SVMs and GBMs appeared to slowly approach the true values with larger sample sizes (results not shown) though some bias still existed at $n = 10\,000$. PSA(II) performed poorly throughout, due to the incorrect specification of the outcome model when conditional on the treatments only as main terms, and did not converge with larger sample sizes. Note that PSA(II) performed similarly to the closely related adjusted G-Computation with treatment main terms. For the second regimen, TMLE(I) was essentially unbiased but often had more variance than IPTW and TMLE(II).

Table 2. Monte Carlo means and standard errors over 1 000 draws for different causal estimators that utilize the generalized propensity score. The true value for regimen 1 is $\mathbb{E}(Y^1) = 0.61$ and the true value for regimen 2 is $\mathbb{E}(Y^2) = 0.57$. SVM: Support Vector Machine; GBM: Generalized Boosted Model; IPTW: Inverse Probability of Treatment Weighting; PSA: Propensity Score Adjustment; TMLE: Targeted Maximum Likelihood Estimation. Outcome regression models were fit by (I) regimen and (II) treatments as main terms covariates. $Q_n \text{ corr}$ indicates whether the outcome model includes the true treatment-treatment interactions.

	$Q_n \text{ corr}$	$n = 500$		$n = 1\ 000$	
		Reg 1	Reg 2	Reg 1	Reg 2
SVM					
IPTW	N/A	0.63(0.05)	0.61(0.07)	0.63(0.04)	0.61(0.05)
PSA(I)	Y	0.64(0.06)	0.51(0.08)	0.64(0.04)	0.52(0.05)
PSA(II)	N	0.44(0.05)	0.83(0.04)	0.44(0.03)	0.83(0.03)
TMLE(I)	Y	0.62(0.05)	0.58(0.09)	0.62(0.04)	0.58(0.06)
TMLE(II)	N	0.62(0.05)	0.60(0.07)	0.62(0.04)	0.59(0.05)
Softmax Regression					
IPTW	N/A	0.62(0.06)	0.58(0.11)	0.62(0.04)	0.58(0.07)
PSA(I)	Y	0.64(0.05)	0.57(0.07)	0.63(0.04)	0.57(0.05)
PSA(II)	N	0.47(0.04)	0.83(0.04)	0.47(0.03)	0.83(0.03)
TMLE(I)	Y	0.62(0.06)	0.58(0.10)	0.62(0.04)	0.58(0.07)
TMLE(II)	N	0.62(0.06)	0.58(0.10)	0.62(0.04)	0.58(0.07)
GBM					
IPTW	N/A	0.62(0.05)	0.60(0.08)	0.62(0.04)	0.59(0.06)
PSA(I)	Y	0.63(0.06)	0.51(0.10)	0.63(0.04)	0.52(0.06)
PSA(II)	N	0.42(0.04)	0.85(0.04)	0.43(0.03)	0.84(0.03)
TMLE(I)	Y	0.62(0.06)	0.58(0.10)	0.62(0.04)	0.58(0.07)
TMLE(II)	N	0.62(0.05)	0.59(0.08)	0.62(0.05)	0.59(0.06)

We conducted a second simulation study with eight dichotomous treatment variables and a sample size of $n = 500$. In our simulated data, out of the 256 possible regimens, roughly 150 different regimens occurred in each dataset. Some of these regimens were only followed by several subjects, making the corresponding GPSs difficult to estimate. We tested whether removing the observations corresponding to the 20 and 30% least supported regimens affected the causal estimation. Specifically, we did not use these observations in the GPS model fitting but kept them in for the other causal estimation steps. We found that, out of a total of 500 observations, this resulted on average in the removal of

only 30 and 45 observations, respectively, reduced the computational time, and did not change the quality of the estimation. We present the full description and the results of this simulation study in the Supplementary Materials Section 2.

4 Application of the above methods to the MDR-TB Data

The Collaborative Group for Meta-Analysis of Individual Patient Data in Multidrug-Resistant Tuberculosis (IPD-MDRTB)⁵ assembled individual patient data on treatment outcomes from 31 observational studies comprising 9 290 individual pulmonary MDR-TB patients. This dataset contains information on the antimicrobial agents used, the baseline covariates (summarized in Table 3), and clinical outcomes. Patients were observed to take 15 different antimicrobial agents in various combinations. We refer to these different sets of medications as regimens and present the ten most prevalent regimens used in the first row of Table 3. Notably, the most common regimens included 5 or more different antimicrobial agents, while 207 subjects did not take any antimicrobial agent. The antimicrobial agents in the ten most observed regimens are ethambutol (EMB), ethionamide (ETH), ofloxacin (OFX), pyrazinamide (Z), kanamycin (KM), cycloserine (CS), capreomycin (CM), para-aminosalicylic acid (PAS), prothionamide (PTO), streptomycin (SM), and rifabutin (RBT).

A binary outcome was defined as either treatment success (the treatment was completed and cured the infection) or failure (patient still tested culture positive for MDR-TB, died, or defaulted on treatment/were lost to follow-up). After removing the 2.77% of subjects with a missing outcome and the 0.34% with missing baseline information, we were left with a sample size of $n = 9\ 001$ observations taking 1 626 different regimens. The covariate age was divided into six categories (0-24, 25-33, 34-42, 43-52, 53-63, 64-) approximately corresponding to age sextiles and the year of study (defined as the final year of patient treatment) was treated as categorical with 14 values. As observed in Table 3, there are differences across the regimen groups in terms of all covariates. This is evidence of indication bias as medication regimens may be differentially assigned across countries, time periods, and patient disease characteristics.

The objective of this data analysis is to compare the results of the different methodological approaches for the estimation of $E(Y^r)$. We do this for the ten most prevalent regimens in the dataset, corresponding to the first ten regimens in

Table 3. Summary of the baseline and outcome data for the application study in Section 4. p : proportion of subjects following regimen; -ve: Negative; +ve: Positive; SD: Standard deviation; IQR: Interquartile range; TB: Tuberculosis; CXR: Chest X-ray

Regimen	1	2	3	4	5	6	7	8	9	10
	OFX-KM- Z-EMB- ETH	OFX-KM- PTO- ETH-CS	OFX-KM- PTO- CS-PAS	Z-EMB- RBT	OFX-SM- PTO- CS-PAS	None	Z- ETH	OFX-KM- Z- ETH-CS-PAS	OFX- PTO- CS-PAS	OFX-KM- Z-EMB- ETH-CS
n	1514	364	263	237	209	207	178	153	151	137
p	0.16	0.04	0.03	0.03	0.02	0.02	0.02	0.02	0.02	0.01
Covariates										
<i>Year, median(IQR)</i>	2008(0)	2008(0)	2004(2)	1997(7)	2002(2)	2002(7)	2008(0)	2004(0)	2004(2)	2008(0)
<i>Country Income Group</i>										
<i>high, p</i>	0.00	0.01	0.97	0.84	1.00	0.61	0.01	0.92	1.00	0.20
<i>lower middle, p</i>	0.00	0.00	0.03	0.00	0.00	0.03	0.00	0.08	0.00	0.00
<i>upper middle, p</i>	1.00	0.99	0.00	0.16	0.00	0.36	0.99	0.00	0.00	0.80
<i>Age, mean(SD)</i>	37.71 (10.67)	36.23 (11.55)	43.51 (15.72)	44.82 (14.60)	44.82 (15.11)	41.07 (15.66)	36.03 (11.90)	36.41 (11.14)	40.23 (15.32)	36.64 (10.40)
<i>Sex, female, p</i>	0.37	0.36	0.22	0.29	0.24	0.35	0.42	0.12	0.29	0.35
<i>HIV</i>										
<i>+ve, p</i>	0.26	0.25	0.00	0.40	0.00	0.21	0.26	0.02	0.00	0.22
<i>-ve, p</i>	0.43	0.37	0.95	0.49	0.99	0.58	0.34	0.97	0.99	0.54
<i>unknown, p</i>	0.31	0.38	0.05	0.11	0.01	0.21	0.40	0.01	0.01	0.24
<i>Past TB</i>										
<i>yes, p</i>	0.90	0.87	0.75	0.24	0.77	0.53	0.90	0.95	0.74	0.93
<i>no, p</i>	0.07	0.06	0.25	0.74	0.22	0.44	0.06	0.00	0.25	0.05
<i>unknown, p</i>	0.03	0.07	0.00	0.02	0.01	0.03	0.04	0.05	0.01	0.02
<i>Sputum smear status</i>										
<i>+ve, p</i>	0.60	0.67	0.66	0.73	0.83	0.58	0.72	0.85	0.72	0.77
<i>-ve, p</i>	0.30	0.23	0.27	0.17	0.17	0.20	0.19	0.12	0.26	0.19
<i>unknown, p</i>	0.10	0.10	0.07	0.10	0.00	0.22	0.09	0.03	0.02	0.04
<i>Cavities on CXR</i>										
<i>+ve, p</i>	0.55	0.64	0.42	0.22	0.46	0.35	0.64	0.54	0.37	0.57
<i>-ve, p</i>	0.10	0.10	0.57	0.26	0.52	0.09	0.15	0.37	0.62	0.18
<i>unknown, p</i>	0.35	0.26	0.01	0.52	0.02	0.56	0.21	0.09	0.01	0.25
Outcome										
<i>Treatment Success, p</i>	0.46	0.62	0.50	0.14	0.38	0.31	0.54	0.71	0.42	0.56

Table 3. The parameter $\mathbb{E}(Y^r)$ can be interpreted as the proportion of the study population that would have had a successful recovery had all the patients been treated with regimen r . Therefore, larger values of this parameter indicate which regimens may be more beneficially applied on a large scale. Ethics approval was obtained for the reanalysis of this data through the Ethics in Health Research Committee at Université de Montréal (certificate number 17-111-CERES-D).

In order to estimate the GPS with SVMs and Softmax Regression, we removed all of the subjects with regimens only supported by one or two subjects (1 420 subjects). The models were fit using the 7 581 remaining observations. The GPS was then predicted for the entire population of $n=9\ 001$ patients conditional on the covariates in Table 3 and indicators for missing values. GBMs were run using the `twang` package and we selected the combinations of interaction depth, `n.minobsinnode` (minimum observations in each node), and shrinkage parameters that produced the best balance using the KS statistic as explained in McCaffrey *and others*⁹. After obtaining the GPS with these methods, we proceeded with the causal estimation procedures described in Section 3.3 for the estimation of $\mathbb{E}(Y^r)$.

Tables 4 and 5 present the estimates of $\mathbb{E}(Y^r)$ obtained for the ten most frequent regimens. No closed-form approximation of the standard error is available for the multi-treatment version of PSA, and given that the machine learning methods were very computationally intensive, numerical methods like bootstrapping weren't feasible for our implementation. Therefore, the confidence intervals for this method were omitted. The logistic regression outcome model used in TMLE(I) overfit the data (causing the update step to fail) and therefore a LASSO penalty was added to the outcome model with penalty parameter chosen using cross validation with the R package `glmnet`³¹. We used empirical summaries of the weights and GPS (Supplementary Materials Sections 5 and 6) to evaluate whether the positivity assumption may be nearly violated for some subjects. No truncation of the GPS³² was used for the results presented, though we conducted a sensitivity analysis where 20% truncation was used to remove the smallest values of the GPS. Numerical results of the sensitivity analyses are presented in the Supplementary Materials Section 6 and discussed below.

Table 4. Estimates of the probability of treatment success along with the confidence intervals under regimens 1-5 for the MDR-TB application in Section 4. SVM: Support Vector Machine; GBM: Generalized Boosted Model; IPTW: Inverse Probability of Treatment Weighting; PSA: Propensity Score Adjustment; TMLE: Targeted Maximum Likelihood Estimation. Outcome regression models were fit (I) by regimen and (II) with treatments as main terms covariates.

Regimen	1	2	3	4	5
	OFX-KM- Z-EMB- ETH	OFX-KM- Z- ETH-CS	OFX-KM- PTO- CS-PAS	Z-EMB- RBT	OFX-SM- PTO- CS-PAS
SVM					
IPTW	0.46 (0.44,0.49)	0.71 (0.62,0.80)	0.59 (0.47,0.70)	0.27 (0.09,0.45)	0.32 (0.17,0.46)
PSA(I)	0.44	0.67	0.63	0.32	0.55
PSA(II)	0.66	0.69	0.64	0.42	0.68
TMLE(I)	0.61 (0.60,0.61)	0.78 (0.76,0.80)	0.63 (0.60,0.65)	0.54 (0.52,0.57)	0.31 (0.28,0.34)
TMLE(II)	0.49 (0.48,0.50)	0.69 (0.68,0.70)	0.60 (0.58,0.63)	0.34 (0.31,0.36)	0.37 (0.35,0.38)
Softmax					
Regression					
IPTW	0.46 (0.43,0.49)	0.65 (0.59,0.70)	0.56 (0.49,0.64)	0.27 (0.18,0.36)	0.37 (0.29,0.44)
PSA(I)	0.38	0.63	0.55	0.22	0.45
PSA(II)	0.56	0.64	0.59	0.36	0.62
TMLE(I)	0.60 (0.59,0.62)	0.65 (0.62,0.67)	0.61 (0.59,0.64)	0.57 (0.54,0.60)	0.37 (0.35,0.39)
TMLE(II)	0.48 (0.47,0.50)	0.64 (0.62,0.67)	0.59 (0.57,0.62)	0.26 (0.22,0.30)	0.45 (0.43,0.48)
GBM					
IPTW	0.55 (0.39,0.72)	0.81 (0.64,0.98)	0.59 (0.47,0.70)	0.25 (0.11,0.39)	0.27 (0.01,0.52)
PSA(I)	0.43	0.68	0.63	0.35	0.55
PSA(II)	0.65	0.68	0.64	0.37	0.66
TMLE(I)	0.63 (0.58,0.68)	0.83 (0.79,0.87)	0.60 (0.54,0.67)	0.54 (0.51,0.56)	0.27 (0.22,0.32)
TMLE(II)	0.55 (0.50,0.60)	0.77 (0.76,0.79)	0.57 (0.50,0.64)	0.34 (0.30,0.37)	0.30 (0.28,0.32)

Table 5. Estimates of the probability of treatment success along with the confidence intervals under regimens 6-10 for the MDR-TB application in Section 4. SVM: Support Vector Machine; GBM: Generalized Boosted Model; IPTW: Inverse Probability of Treatment Weighting; PSA: Propensity Score Adjustment; TMLE: Targeted Maximum Likelihood Estimation. Outcome regression models were fit (I) by regimen and (II) with treatments as main terms covariates.

Regimen	6 None	7 OFX-KM- Z- ETH	8 OFX-CM- Z- ETH-CS-PAS	9 OFX- PTO- CS-PAS	10 OFX-KM- Z-EMB- ETH-CS
SVM					
IPTW	0.20 (0.08,0.31)	0.56 (0.48,0.64)	0.67 (0.55,0.80)	0.57 (0.37,0.77)	0.56 (0.47,0.64)
PSA(I)	0.29	0.59	0.61	0.56	0.57
PSA(II)	0.38	0.63	0.61	0.58	0.66
TMLE(I)	0.21 (0.18,0.23)	0.58 (0.56,0.60)	0.67 (0.65,0.69)	0.62 (0.58,0.66)	0.61 (0.58,0.63)
TMLE(II)	0.24 (0.21,0.27)	0.58 (0.56,0.60)	0.60 (0.58,0.62)	0.58 (0.55,0.61)	0.57 (0.54,0.60)
Softmax					
Regression					
IPTW	0.31 (0.24,0.38)	0.56 (0.48,0.64)	0.69 (0.61,0.78)	0.45 (0.35,0.54)	0.56 (0.48,0.65)
PSA(I)	0.37	0.55	0.56	0.46	0.54
PSA(II)	0.38	0.56	0.59	0.50	0.65
TMLE(I)	0.25 (0.22,0.28)	0.58 (0.55,0.61)	0.68 (0.66,0.70)	0.55 (0.52,0.58)	0.60 (0.57,0.64)
TMLE(II)	0.35 (0.29,0.41)	0.56 (0.53,0.60)	0.62 (0.60,0.64)	0.49 (0.47,0.52)	0.56 (0.52,0.61)
GBM					
IPTW	0.24 (0.17,0.32)	0.70 (0.41,0.98)	0.75 (0.65,0.83)	0.56 (0.25,0.86)	0.55 (0.45,0.65)
PSA(I)	0.38	0.60	0.60	0.54	0.57
PSA(II)	0.40	0.62	0.60	0.52	0.65
TMLE(I)	0.25 (0.21,0.28)	0.67 (0.62,0.73)	0.73 (0.70,0.77)	0.62 (0.56,0.67)	0.59 (0.57,0.62)
TMLE(II)	0.26 (0.22,0.31)	0.67 (0.65,0.68)	0.67 (0.63,0.71)	0.58 (0.54,0.61)	0.54 (0.52,0.58)

The point estimates of $\mathbb{E}(Y^r)$ and the confidence intervals in Tables 4 and 5 often varied depending on which method was used to estimate the GPS. The point estimates also sometimes disagreed between causal inference methods using the same GPS vector (e.g. regimens 1 (OFX-KM-Z-EMB-ETH) and 5 (OFX-SM-PTO-CS-PAS)) and to a lesser extent between GPS methods using the same causal inference method. None of the GPS methods consistently produced narrow confidence intervals for TMLE or IPTW. However, TMLE was often found to have narrower confidence intervals than IPTW. GPS truncation resulted in at most small changes in the point estimates though very small values of the GPS were observed, suggesting possible near-positivity violations.

Table 6 presents the top 5 most beneficial regimens based on the estimates of $\mathbb{E}(Y^r)$. Regimens 2 (OFX-KM-Z-ETH-CS) and 8 (OFX-CM-Z-ETH-CS-PAS) were often classified in the top 2 and were in the top 5 of all methods except for PSA(II) with SVMs and GBMs. Regimens 3 (OFX-KM-PTO-CS-PAS), 7 (OFX-KM-Z-ETH), and 10 (OFX-KM-Z-EMB-ETH-CS) were also often ranked in the top 5. This would suggest the superior effectiveness of these treatment combinations among the regimens investigated.

World Health Organization (WHO) guidelines^{4,33} suggest that MDR-TB regimens include a fluoroquinolone (such as OFX) and an injectable agent (such as KM, SM or CM). No treatment (included as a benchmark despite questionable clinical interest) and regimen 4 (Z-EMB-RBT) performed the worst overall and follow neither of these guidelines. Regimen 9 (OFX-PTO-CS-PAS), which also performed poorly, also lacked an injectable agent.

WHO guidelines also point to the importance of the number of drugs in the regimen, suggesting five or more that have certain or almost certain effectiveness⁴. Previous studies have suggested that a majority of MDR-TB patients are resistant to EMB and Z in many settings^{34,35}. When excluding EMB and Z, of the regimens evaluated here, regimens 3, 5, and 8 had five remaining drugs, though only regimens 3 and 8 were found to be among the most effective. Regimen 5 was identical to regimen 3 except that it replaced KM by SM, for which resistance is also commonly seen among MDR-TB isolates. Regimens 2, 9, and 10 had four remaining drugs. Regimens 2 and 10 both included an injectable (KM) and were identical except that regimen 10 also included EMB. Interestingly, regimen 2 was found to perform the best among the regimens evaluated while 10 was found to be less effective. These results point to the potential importance of the inclusion

of KM in a regimen. While we estimated the expected mean of the potential outcome under ten regimens, future applications may use marginal structural models (modeling of the expected potential outcomes conditional on treatments) and a broader range of regimens to estimate the contributions of each individual treatment and treatment interaction on the outcome. In the discussion, we point out some limitations of the simplified analysis in the current paper, which limits the interpretability of the results.

Table 6. Ranking of the top 5 medication regimens estimated by each method in terms of the estimated population recovery rate of MDR-TB treatment success. *Reg 1:* OFX-KM-Z-EMB-ETH; *Reg 2:* OFX-KM-Z-ETH-CS; *Reg 3:* OFX-KM-PTO-CS-PAS; *Reg 4:* Z-EMB-RBT; *Reg 5:* OFX-SM-PTO-CS-PAS; *Reg 6:* None; *Reg 7:* OFX-KM-Z-ETH; *Reg 8:* OFX-CM-Z-ETH-CS-PAS; *Reg 9:* OFX-PTO-CS-PAS; *Reg 10:* OFX-KM-Z-EMB-ETH-CS; SVM: Support Vector Machine; GBM: Generalized Boosted Model; IPTW: Inverse Probability of Treatment Weighting; PSA: Propensity Score Adjustment; TMLE: Targeted Maximum Likelihood Estimation. Outcome regression models were fit (I) by regimen and (II) with treatments as main terms covariates.

Causal Estimation Methods	IPTW	TMLE(I)	TMLE(II)	PSA(I)	PSA(II)
SVM	<i>Reg 2</i>	<i>Reg 2</i>	<i>Reg 2</i>	<i>Reg 2</i>	<i>Reg 2</i>
	<i>Reg 8</i>	<i>Reg 8</i>	<i>Reg 8</i>	<i>Reg 3</i>	<i>Reg 5</i>
	<i>Reg 3</i>	<i>Reg 3</i>	<i>Reg 3</i>	<i>Reg 8</i>	<i>Reg 10</i>
	<i>Reg 9</i>	<i>Reg 9</i>	<i>Reg 7</i>	<i>Reg 7</i>	<i>Reg 1</i>
	<i>Reg 10</i>	<i>Reg 10</i>	<i>Reg 9</i>	<i>Reg 10</i>	<i>Reg 3</i>
Softmax Regression	<i>Reg 8</i>	<i>Reg 8</i>	<i>Reg 2</i>	<i>Reg 2</i>	<i>Reg 10</i>
	<i>Reg 2</i>	<i>Reg 2</i>	<i>Reg 8</i>	<i>Reg 8</i>	<i>Reg 2</i>
	<i>Reg 10</i>	<i>Reg 3</i>	<i>Reg 3</i>	<i>Reg 3</i>	<i>Reg 5</i>
	<i>Reg 7</i>	<i>Reg 10</i>	<i>Reg 10</i>	<i>Reg 7</i>	<i>Reg 8</i>
	<i>Reg 3</i>	<i>Reg 1</i>	<i>Reg 7</i>	<i>Reg 10</i>	<i>Reg 3</i>
GBM	<i>Reg 2</i>	<i>Reg 2</i>	<i>Reg 2</i>	<i>Reg 2</i>	<i>Reg 2</i>
	<i>Reg 8</i>	<i>Reg 8</i>	<i>Reg 8</i>	<i>Reg 3</i>	<i>Reg 5</i>
	<i>Reg 7</i>	<i>Reg 7</i>	<i>Reg 7</i>	<i>Reg 7</i>	<i>Reg 10</i>
	<i>Reg 3</i>	<i>Reg 1</i>	<i>Reg 9</i>	<i>Reg 8</i>	<i>Reg 1</i>
	<i>Reg 9</i>	<i>Reg 9</i>	<i>Reg 3</i>	<i>Reg 10</i>	<i>Reg 3</i>

5 Discussion

In this paper, we investigated the causal estimation of multiple concurrent medications as motivated by the clinical question of how best to treat patients with MDR-TB. The topic of polypharmacy (resulting in potential overmedication and dangerous medication interactions) is gaining in importance in the medical literature. In particular, polypharmacy is highly prevalent in the elderly (ages ≥ 65), an important and growing population³⁶ leading to potential adverse drug reactions³⁷. For example, multiple cardiovascular medications, taken by more than 50% of elderly people, have been shown to be associated with an increased risk of acute kidney disorders³⁸. Given the toxicity of second-line anti-tuberculosis drugs, the analysis of polypharmacy is particularly relevant for treating MDR-TB cases.

In order to address estimation in this challenging scenario, we defined a treatment “regimen” as each unique combination of medications and used three methods to estimate the GPS, or the probability of receiving a specific regimen. One weakness of this GPS approach is that it does not directly allow for information to be shared between different regimens that contain one or more of the same medications. In a Monte Carlo simulation study, we showed that missing treatment interactions in the outcome model can lead to bias in the estimation of both PSA and G-Computation. In real world applications, it might therefore be difficult to correctly specify these models. However, due to its double robustness property, TMLE was found to produce unbiased point estimates even when the outcome model was incorrectly specified. Further investigations could involve the implementation of TMLE with a non-parametric method used for the outcome model as well, which might add additional robustness to the estimation³⁹.

In the application, we estimated the probability of treatment success for the 10 most prevalent medication combinations in the MDR-TB dataset. We chose to estimate the most prevalent medications because they may be of greatest clinical interest and also have the greatest amount of data support (i.e. number of patients following the regimens) which allowed for better estimation. An interesting question for future research would involve empirically identifying which regimens have sufficient data support. One may also integrate existing methods to data-adaptively select covariates in the GPS for a given regimen⁴⁰.

The different methods often agreed on the preferred MDR-TB regimens but produced sometimes differing estimates of probabilities of treatment success. Closed-form confidence intervals are not available for PSA with multiple regimens

and we were unable to use a numerical approach to approximate them given the computational complexity of the GPS estimation. Previous investigations of this data source⁵ used regression analyses to estimate the associations between each treatment and outcome separately, ignoring other treatments. Associations between the number of treatments and duration of treatment were also investigated. In contrast to the previous approach, our general approach considers the joint effect of treatments. TMLE also has the advantage of being doubly robust and therefore consistent when either the GPS or the outcome model is correctly specified. Since the dataset consists of the fusion of multiple observational studies, a more appropriate application of these methods would formally consider the heterogeneity between studies in the point estimation (e.g. using a random effects outcome model by study) and standard error estimation⁴¹ and account for selection bias as different populations were observed to take different regimens of antimicrobial agents. Our analysis also did not consider known drug resistance in the analysis, which may affect treatment decisions and outcomes, nor did we address the extrapolation required to synthesize evidence when certain regimens are only observed in select time periods. Ongoing analyses more appropriately address these issues and strong clinical conclusions about medication or regimen effectiveness are beyond the scope of this article.

Because of the large number of regimens, the GPS model may sometimes predict very small probabilities for some regimens. This creates well-known stability problems for methods that weight by the inverse of the GPS. We addressed this problem by using formulations of IPTW and TMLE that use the inverse GPS as a weight in a regression. Alternative approaches (results not shown) were sometimes highly biased in the simulation study. The robustness of the regression approach is likely due to the dampening of the residuals in the weighted regression step. TMLE and IPTW often benefit from GPS truncation as a bias-variance trade-off and data-adaptive approaches have been recently proposed⁴². However, small values of the GPS may also indicate true positivity violations and the nonexistence of the parameter of interest. Very small values of the GPS could be investigated to identify patients who were truly ineligible for a given treatment due to clinical or demographic features. Related to the simplifications mentioned above, we did not consider this possibility.

An alternative approach that we considered but did not take in this paper (that addresses the mentioned limitation) involves treating the regimens not

as categorical, but as a multivariate binary variable, with each component indicating whether a subject was on that specific medication. One could then attempt to use multivariate regression modeling⁴³ for the GPS that allows for some correlation between the treatments. Optimal Classifier Chains⁴⁴, or simpler regression approaches that otherwise allow for dependencies between the usage of different treatments are potential approaches.

It is clear from the medical literature that the estimation of the effects of multiple concurrent medications is an important topic but standard methods are limited. Given the complexity of the problem, we hope that this paper encourages additional focus on these methodological issues.

Code and Supplementary Materials

The code for the simulation study data generation and analysis is available at <https://github.com/arman817/Simulation-Codes-for-Causal-Inference-for-polypharmacy>. The extended simulation and application results are included in a Supplementary Materials file, available online.

Acknowledgements

The data and context were provided by the Collaborative Group for Meta-Analysis of Individual Patient Data in Multidrug-Resistant Tuberculosis. The authors gratefully acknowledge Matthew Cefalu's (Rand Corporation) recommendations for the implementation of `twang`.

Funding

This work was supported by the Canadian Institutes of Health Research (CIHR) [Project Grant 378067 to MES and AB]. MES is also funded by CIHR [New Investigators Salary Award] and the National Sciences and Engineering Council of Canada [Discovery Grant with Accelerator Supplement]. NRG is funded in part by the National Institutes of Health (NIH) [K24 award, K24AI114444]. The funders had no role in the study design, data collection and analysis, decision to publish, or preparation of the manuscript. The corresponding author had full access to all data in the study and had final responsibility for the decision to submit for publication. The findings and conclusions in this report are

those of the authors and do not necessarily represent the official position of the funding agencies.

References

1. Maher RL, Hanlon J and Hajjar ER. Clinical consequences of polypharmacy in elderly. *Expert Opinion on Drug Safety* 2014; 13(1): 57–65.
2. World Health Organization. *Global tuberculosis report 2017*. World Health Organization, 2017.
3. Millard J, Ugarte-Gil C and Moore DA. Multidrug resistant tuberculosis. *The BMJ* 2015; 350: h882.
4. World Health Organization. *WHO Treatment Guidelines for Drug-resistant Tuberculosis 2016 Update*. World Health Organization, 2016.
5. Ahuja SD, Ashkin D, Avendano M et al. Multidrug resistant pulmonary tuberculosis treatment regimens and patient outcomes: an individual patient data meta-analysis of 9,153 patients. *PLoS Medicine* 2012; 9(8): e1001300.
6. Imbens GW. The role of the propensity score in estimating dose-response functions. *Biometrika* 2000; 87(3): 706–710.
7. Imai K and Van Dyk DA. Causal inference with general treatment regimes: Generalizing the propensity score. *Journal of the American Statistical Association* 2004; 99(467): 854–866.
8. Lopez MJ and Gutman R. Estimation of causal effects with multiple treatments: a review and new ideas. *arXiv preprint arXiv:170105132* 2017; .
9. McCaffrey DF, Griffin BA, Almirall D et al. A tutorial on propensity score estimation for multiple treatments using generalized boosted models. *Statistics in Medicine* 2013; 32(19): 3388–3414.
10. Horvitz DG and Thompson DJ. A generalization of sampling without replacement from a finite universe. *Journal of the American Statistical Association* 1952; 47(260): 663–685.
11. Vansteelandt S and Daniel RM. On regression adjustment for the propensity score. *Statistics in Medicine* 2014; 33(23): 4053–4072.
12. Scharfstein DO, Rotnitzky A and Robins JM. Adjusting for nonignorable dropout using semiparametric nonresponse models (with discussion and rejoinder). *Journal of the American Statistical Association* ; 94(448): 1121–1146.
13. Van der Laan MJ and Rubin D. Targeted maximum likelihood learning. *The International Journal of Biostatistics* 2006; 2(1).
14. Robins J. 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* 1986; 7(9-12): 1393–1512.
15. Rosenbaum PR and Rubin DB. The central role of the propensity score in observational studies for causal effects. *Biometrika* 1983; 70(1): 41–55.
 16. Hastie T, Tibshirani R and Friedman J. *The Elements of Statistical Learning*. 2 ed. Springer, New York, NY, 2009.
 17. Ahuja Y and Yadav SK. Multi-class classification and support vector machine. *Global Journal of Computer Science and Technology* 2012; 12(11): 14–20.
 18. Cortes C and Vapnik V. Support-vector networks. *Machine Learning* 1995; 20(3): 273–297. DOI:10.1007/BF00994018. URL <https://doi.org/10.1007/BF00994018>.
 19. Meyer D, Dimitriadou E, Hornik K et al. Package e1071, 2017.
 20. Hsu CW and Lin CJ. A comparison of methods for multiclass support vector machines. *IEEE Transactions on Neural Networks* 2002; 13(2): 415–425. DOI:10.1109/72.991427.
 21. Karatzoglou A, Meyer D and Hornik K. Support vector machines in R. *Journal of Statistical Software, Volume* 2006; : 1–28.
 22. Wu TF, Lin CJ and Weng RC. Probability estimates for multi-class classification by pairwise coupling. *Journal of Machine Learning Research* 2004; 5: 975–1005.
 23. Hosmer Jr DW, Lemeshow S and Sturdivant RX. *Applied logistic regression, Wiley Series in Probability and Statistics*, volume 3. John Wiley & Sons, 2013.
 24. Ding X. *softmaxreg: Training Multi-Layer Neural Network for Softmax Regression and Classification*, 2016. URL <https://CRAN.R-project.org/package=softmaxreg>. R package version 1.2.
 25. Ridgeway G, Mccaffrey D, Morral A et al. Twang: Toolkit for weighting and analysis of nonequivalent groups. R package version 1.0–1. In *R Foundation for Statistical Computing, 2006*.
 26. Snowden JM, Rose S and Mortimer KM. Implementation of g-computation on a simulated data set: demonstration of a causal inference technique. *American Journal of Epidemiology* 2011; 173(7): 731–738.
 27. Zeileis A. Object-oriented computation of sandwich estimators. *Journal of Statistical Software* 2006; 16(9).
 28. Van der Laan MJ and Rose S. *Targeted learning: causal inference for observational and experimental data*. Springer Science & Business Media, 2011.

29. Tsiatis A. *Semiparametric theory and missing data*. Springer Science & Business Media, 2007.
30. Kennedy EH. Semiparametric theory and empirical processes in causal inference. Technical report, Cornell University Library, 2016. URL <https://arxiv.org/abs/1510.04740>.
31. Friedman J, Hastie T and Tibshirani R. Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software* 2010; 33(1): 1.
32. Cole SR and Hernán MA. Constructing inverse probability weights for marginal structural models. *American Journal of Epidemiology* 2008; 168(6): 656–664.
33. World Health Organization. *Guidelines for the programmatic management of drug-resistant tuberculosis*. World Health Organization, Geneva, 2006.
34. Munir S, Mahmood N, Shahid S et al. Molecular detection of isoniazid, rifampin and ethambutol resistance to *m. tuberculosis* and *m. bovis* in multidrug resistant tuberculosis (mdr-tb) patients in pakistan. *Microbial Pathogenesis* 2017; 110: 262–274.
35. Allana S, Shashkina E, Mathema B et al. *pnca* gene mutations associated with pyrazinamide resistance in drug-resistant tuberculosis, south africa and georgia. *Emerging Infectious Diseases* 2017; 23(3): 491–495.
36. US Department of Health and Human Services. A profile of older americans: 2016. Technical report, 2016.
37. Qato DM, Wilder J, Schumm LP et al. Changes in prescription and over-the-counter medication and dietary supplement use among older adults in the united states, 2005 vs 2011. *JAMA Internal Medicine* 2016; 176(4): 473–482.
38. Chao CT, Tsai HB, Wu CY et al. Cumulative cardiovascular polypharmacy is associated with the risk of acute kidney injury in elderly patients. *Medicine* 2015; 94(31): e1251.
39. Benkeser D, Carone M, Van der Laan MJ et al. Doubly robust nonparametric inference on the average treatment effect. *Biometrika* 2017; 104(4): 863–880. DOI:10.1093/biomet/asx053.
40. Gruber S and van der Laan MJ. *C-TMLE of an Additive Point Treatment Effect*. New York, NY: Springer New York, 2011. pp. 301–321.
41. Schnitzer ME, Van der Laan MJ, Moodie EEM et al. Effect of breastfeeding on gastrointestinal infection in infants: a targeted maximum likelihood approach for clustered longitudinal data. *The Annals of Applied Statistics* 2014; 8(2): 703–725.
42. Ju C, Schwab J and van der Laan MJ. On adaptive propensity score

- truncation in causal inference. *Statistical Methods in Medical Research* 2018; 0(0): in press.
43. Ip S and Xue J. A multivariate regression view of multi-label classification. Technical report, University College London, 2015. URL <http://www.ucl.ac.uk/~zcapg66/CSML/report.pdf>.
 44. Cheng W, Hüllermeier E and Dembczynski KJ. Bayes optimal multilabel classification via probabilistic classifier chains. In *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*. pp. 279–286.