


Confounding adjustment methods for multi-level treatment comparisons under lack of positivity and unknown model specification

S. Arona Diop^a, Thierry Duchesne^a, Steven G. Cumming^b, Awa Diop^c and Denis Talbot ^c

^aDépartement de mathématiques et de statistique, Université Laval, Québec, Canada; ^bDépartement des sciences du bois et de la forêt, Université Laval, Québec, Canada; ^cDépartement de médecine sociale et préventive, Université Laval, Québec, Canada

ABSTRACT

Imbalances in covariates between treatment groups are frequent in observational studies and can lead to biased comparisons. Various adjustment methods can be employed to correct these biases in the context of multi-level treatments (> 2). Analytical challenges, such as positivity violations and incorrect model specification due to unknown functional relationships between covariates and treatment or outcome, may affect their ability to yield unbiased results. Such challenges were expected in a comparison of fire-suppression interventions for preventing fire growth. We identified the overlap weights, augmented overlap weights, bias-corrected matching and targeted maximum likelihood as methods with the best potential to address those challenges. A simple variance estimator for the overlap weight estimators that can naturally be combined with machine learning is proposed. In a simulation study, we investigated the performance of these methods as well as those of simpler alternatives. Adjustment methods that included an outcome modeling component performed better than those that focused on the treatment mechanism in our simulations. Additionally, machine learning implementation was observed to efficiently compensate for the unknown model specification for the former methods, but not the latter. Based on these results, we compared the effectiveness of fire-suppression interventions using the augmented overlap weight estimator.

ARTICLE HISTORY

Received 20 August 2020
Accepted 29 March 2021


KEYWORDS

Multi-level treatment;
machine learning;
confounding adjustment;
plasmode simulation;
simulation

1. Introduction

Many empirical studies seek to evaluate the effect of a treatment or intervention. This evaluation can be attempted using randomized or observational experiments. In the former, pre-randomization characteristics are expected to be similar across treatment groups. As such, outcome differences can be causally attributed to the treatment. However, randomized trials are often difficult to realize because they may be unethical, impractical,

CONTACT Denis Talbot  denis.talbot@fmed.ulaval.ca

 Supplemental data for this article can be accessed here. <https://doi.org/10.1080/02664763.2021.1911966>

© 2021 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group
This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives License (<http://creativecommons.org/licenses/by-nc-nd/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited, and is not altered, transformed, or built upon in any way.

or untimely [12]. Thus, relying on observational experiments is often necessary. Unfortunately, imbalances in covariates between treatment groups are frequent in observational studies and can lead to biased treatment comparisons. This major challenge is known as confounding, wherein differences in outcomes between treatment groups are at least partly attributable to systematic differences in baseline covariates.

The motivation for the current study was the comparison of the effect of multiple fire-suppression interventions on fire growth in Alberta, Canada. In a previous attempt at tackling this question using observational data, a reversed association was observed wherein more aggressive interventions were associated with greater fire growth [37]. We believed two main reasons could explain this unexpected association. First, important confounding by indication appeared plausible, since more aggressive interventions would be preferred for wildfires that are expected to grow quickly. Adjusting for further potential confounders may help reducing bias, but how best to adjust for these confounders was unclear. We also hypothesized that the positivity assumption may be violated in these data. Positivity refers to the fact that all units should have a positive probability of receiving each possible treatment and is required for many adjustment methods to yield unbiased estimators. In our context, it seemed possible, for example, that some fires may be too large when initially discovered to consider less aggressive interventions, thus violating the positivity assumption.

Various methods can be used to correct confounding bias. The performance of these methods have been compared in multiple simulation studies in the case of a binary treatment [4–6,15,20,21,26,28] and for performing pairwise comparisons between multi-level treatments [8,22,25,35,42,43]. In summary, methods that rely on binary propensity score for multi-level treatments comparisons, focusing on comparing treatment levels two at a time, can yield biased and nontransitive comparisons and are not recommended [25,42]. Stratification was susceptible to produce results with important residual bias [4–6,26]. Standardization is generally the best performing method in ideal situations, in terms of bias and mean-squared error (MSE) [4,15,20,21,26]. However, when the model used for standardization is incorrectly specified, the estimator is biased [8,20,21,26]. Weighting methods were susceptible to be biased and have a large variance if the model used for constructing the weights was misspecified [15,20,42] or under positivity violation [22,25,28,35]. Trimming, which removes from the analysis observations in areas where there is a lack of positivity, has been observed to improve the performance of multiple methods in presence of positivity violations [15,22]. The augmented inverse probability weighting estimator has been observed to have mixed performances. This estimator performed competitively to standardization in ideal situations and retained similar performances under model misspecifications [26], but it performed poorly under positivity violations [15,21]. The targeted maximum likelihood estimator (TMLE) was observed to have good performances under ideal situations and under model misspecifications, in addition to being more robust to positivity violations [20]. Similar results were observed concerning bias-corrected matching [15,20,21]. The use of machine learning algorithms was observed to be beneficial for preventing the bias attributable to model misspecifications for TMLE, BCM and standardization, but not necessarily for weighting [20].

From a theoretical perspective, methods that combine treatment and outcome modeling are generally expected to be \sqrt{n} -consistent as long as each component is consistently

estimated at a rate $\geq n^{1/4}$ (see for example [17] and references therein). This allows using a greater array of machine learning methods without sacrificing desirable asymptotic properties. This is unlike methods that rely on modeling only the treatment or the outcome which typically require \sqrt{n} -consistency of their model to be \sqrt{n} -consistent.

Weighting estimators that are robust to positivity violations have also been developed in recent years [21–23,43]. In particular, [23] consider a general class of weighting estimators for binary treatments, that notably include as specific cases standard weighting and weighting with trimmed weights. They theoretically derived, under certain conditions, the lowest variance estimator in this class, which they called the overlap weights. An extension of these overlap weights to multi-level treatments has also been developed [22]. The overlap weights were observed to outperform the usual weighting estimator, with or without trimming, as well as matching estimators, both in terms of bias and MSE [22,28]. An outcome regression augmented overlap weighting estimator can be constructed in order to improve the robustness to model misspecification [22,28].

Overall, a few methods thus stand out according to their performance in terms of bias, MSE, and their robustness to model misspecifications or positivity violations: bias-corrected matching, TMLE and overlap weights. While bias-corrected matching and TMLE have been compared in simulation studies with binary treatments, they have never been compared together in scenarios with multi-level treatment or against the overlap weights. Which of these methods yield the best estimator of pairwise average treatment effects, including in scenarios with model misspecifications or positivity violations, is thus unknown. To inform our analysis of the wildfire data, we investigated the performance of TMLE, bias-corrected matching and overlap weights for performing pairwise comparisons between multi-level treatments, notably under lack of positivity and unknown model specification.

The remainder of this article is structured as follows. In the next section, the notation used throughout the article is first introduced. Then, the methods being compared are formally presented, as well as the causal assumptions on which they rely. We also propose some methodological developments concerning the overlap weight estimators, including a convenient asymptotic variance estimator. In Section 3, we present the simulation study we have conducted to compare bias-corrected matching, TMLE, overlap weights, and augmented overlap weights. Standardization, inverse probability weighting and matching were also included as benchmark comparators. We further examine the potential of machine learning implementations of these methods to deal with the unknown model specification. Both a fully synthetic Monte Carlo simulation, and a plasmode simulation are presented. Plasmode simulations combine real and synthetic data to produce more realistic simulation scenarios [10]. Section 4 presents the analysis of the wildfire data, which is informed by the results of our simulation study. We conclude with a discussion in Section 5.

2. Methodology

Let Y be the outcome of interest, $T \in \mathcal{T} = \{1, 2, \dots, k\}$ be the multi-level treatment, and $\mathbf{X} = \{X_1, \dots, X_p\} \in \mathcal{X}$ be a set of covariates, where \mathcal{T} and \mathcal{X} are the support of variables T and \mathbf{X} , respectively. We assume that a sample of independent observations $\{\mathbf{X}_i, T_i, Y_i\}$, with $i = 1, \dots, n$, is drawn from a given population. We also denote by $Y_i(t)$ the potential

outcome of observation i under treatment level t , that is, the outcome that would have been observed for observation i had the treatment taken level t , potentially contrary to the fact. Using this notation, the population average pairwise treatment effect comparing treatment levels t and t' is $\tau_{t,t'} = \mathbb{E}[Y(t) - Y(t')]$.

The fundamental problem of causal inference is that it is only possible to observe one potential outcome for each observation, the one corresponding to the factual treatment level. Since the other potential outcomes are missing, the causal effect $\tau_{t,t'}$ cannot be estimated from the observed data without making some assumptions. Henceforth, we make the following usual assumptions (see, for example, [12, Chapter 3]): (1) $Y(t) \perp\!\!\!\perp T | \mathbf{X}$ for all $t \in \mathcal{T}$, (2) $0 < P(T = t | \mathbf{X}) < 1$ for all $t \in \mathcal{T}$ and $\mathbf{X} \in \mathcal{X}$ and (3) $T_i = t \Rightarrow Y_i = Y_i(t)$. The first assumption is exchangeability and indicates that treatment allocation is independent of the potential outcomes conditional on measured covariates. Informally, this means there are no factors that would simultaneously explain T and Y beyond what may be explained by \mathbf{X} . Assumption (2) is the positivity assumption that entails that each observation had a strictly positive probability of being assigned to any treatment level. The third is the consistency assumption and indicates that the observed outcome for observation i is equal to its potential outcome under its factual treatment level.

2.1. Standardization

A first possible estimator of $\tau_{t,t'}$ consists of computing a standardized mean difference of the outcome between treatment groups t and t' (see, for example, [12, Chapter 13]):

$$\hat{\tau}_{t,t'}^{stan} = \frac{1}{n} \sum_{i=1}^n \left[\hat{\mathbb{E}}(Y | T = t, \mathbf{X}_i) - \hat{\mathbb{E}}(Y | T = t', \mathbf{X}_i) \right],$$

where $\hat{\mathbb{E}}(Y | T = t, \mathbf{X}_i)$ and $\hat{\mathbb{E}}(Y | T = t', \mathbf{X}_i)$ are the estimated expectations of the outcome under treatment level t and t' , respectively, and covariates \mathbf{X}_i . These estimated expectations can be obtained using a parametric regression model of Y conditional on T and \mathbf{X} , such as a linear regression if Y is continuous or a logistic regression if Y is binary. More sophisticated methods, including machine-learning algorithms, could also be employed. Although a variance estimator for $\hat{\tau}_{t,t'}^{stan}$ that accounts for the estimation of treatment model parameters exists [44], it can be computationally intensive since it involves a double sum over the observations' index. Alternatively, inferences can be produced using the nonparametric bootstrap, which also allows accounting for parameters estimation.

2.2. Inverse probability weighting

Inverse probability weighting estimators of $\tau_{t,t'}$ are inspired by the method proposed by Hovitz and Thompson [14]. Robins [32] introduced inverse probability weighting estimators of the parameters of longitudinal marginal structural models, which include as a special case the following estimator of $\tau_{t,t'}$:

$$\hat{\tau}_{t,t'}^{IPW} = \frac{1}{n} \sum_{i=1}^n \left[\frac{I(T_i = t)Y_i - I(T_i = t')Y_i}{\hat{P}(T = t_i | \mathbf{X}_i)} \right],$$

where $I(\cdot)$ denotes the usual indicator function and $\hat{P}(T = t_i | \mathbf{X}_i)$ is the estimated probability that the treatment takes level t_i conditional on \mathbf{X}_i . $\hat{P}(T = t_i | \mathbf{X}_i)$ can, for example, be obtained from a multinomial logistic regression. A conservative asymptotic variance estimator for $\hat{\tau}_{t,t'}^{IPW}$ is conveniently obtained using a sandwich estimator, treating $\hat{P}(T = t_i | \mathbf{X}_i)$ as known [26,33].

2.3. Matching and bias-corrected matching

Scotina et al. [35] present a matching estimator as well as a bias-corrected matching estimator based on the work of Abadie and Imbens [1,2]. This matching estimator first imputes the missing potential outcomes of each observation by the observed outcomes of observations from the other treatment group. More precisely, the potential outcome of observation i under treatment t , $Y_i(t)$, for $t \neq t_i$, is imputed with the average of the observed outcome of the m observations from treatment group t whose covariates' values are closest to \mathbf{X}_i according to some distance metric. The potential outcome $Y_i(t_i)$ is deemed observed ($Y_i(t_i) = Y_i$). This matching procedure is performed with replacement, allowing each observation to be used multiple times for imputing potential outcomes of different observations. Comparisons between treatment groups are then performed directly on the (observed/imputed) potential outcomes.

More formally, let A be some definite positive matrix and $\|x\|_A = (x'Ax)^{1/2}$ define the distance metric. Let M_i^t denote the set of the m observations 'closest' to unit i in treatment group $t \neq t_i$, that is, such that $\sum_{j \in M_i^t} \|X_j - X_i\|$ is minimal and $\sum_j I(j \in M_i^t) = m$. Then

$$\hat{Y}_i(t) = \begin{cases} Y_i, & \text{if } t = t_i, \\ \frac{1}{m} \sum_{j \in M_i^t} Y_j, & \text{if } t \neq t_i, \end{cases}$$

and the matching estimator is

$$\hat{\tau}_{t,t'}^{match} = \frac{1}{n} \sum_{i=1}^n [\hat{Y}_i(t) - \hat{Y}_i(t')].$$

Similar matching estimators for multi-level treatments had previously been introduced by [9] and [42].

The bias-corrected matching estimator proposed by [35] uses a regression method for imputing the missing potential outcomes:

$$\hat{Y}_i^{bc}(t) = \begin{cases} Y_i, & \text{if } t = t_i, \\ \frac{1}{m} \sum_{j \in M_i^t} [Y_j + \hat{\mathbb{E}}(Y|T = t, \mathbf{X}_i) - \hat{\mathbb{E}}(Y|T = t, \mathbf{X}_j)], & \text{if } t \neq t_i. \end{cases}$$

The bias-corrected matching estimator is:

$$\hat{\tau}_{t,t'}^{BCM} = \frac{1}{n} \sum_{i=1}^n [\hat{Y}_i^{bc}(t) - \hat{Y}_i^{bc}(t')].$$

The consistency of the matching estimator $\hat{\tau}_{t,t'}^{match}$ and the bias-corrected matching estimator $\hat{\tau}_{t,t'}^{BCM}$ relies on further conditions, in addition to the causal assumptions presented earlier (for more details, see [35]). Asymptotic variance estimators of $\hat{\tau}_{t,t'}^{match}$ and of $\hat{\tau}_{t,t'}^{BCM}$ are provided in [35]. These variance estimators were derived assuming the matching is performed on fixed \mathbf{X} and bias-correction employs a non-parametric series estimator for the outcome expectation. However, the variance estimators were observed to behave well in other circumstances in simulation studies, notably when matching on the logit of the estimated generalized propensity score vector $\hat{P}(T = t|\mathbf{X})$ [35].

2.4. Targeted maximum likelihood

Targeted maximum likelihood estimation is a general methodology introduced by van der Laan and Rubin [40] for constructing doubly-robust semiparametric efficient estimators. Double-robustness entails that the estimator is consistent for $\tau_{t,t'}$ if either the outcome model or the treatment model component of the algorithm described below is consistent, but not necessarily both. Moreover, the TMLE has minimal variance among the class of semiparametric estimators when both models are correctly specified. Luque-Fernandez *et al.* [27] provide an introductory tutorial on TMLE.

An algorithm for obtaining a TMLE, if Y is binary, is as follows [36].

- For $j \in \{t, t'\}$
 - (1) Denote $Q^0(j, \mathbf{X}_i) = \hat{\mathbb{E}}(Y|T = j, \mathbf{X}_i)$
 - (2) Define weights $w_i(j) = \frac{I(T_i=j)}{\hat{P}(T=j|\mathbf{X}_i)}$
 - (3) Run a logistic regression of Y with an intercept, $\text{logit}[Q^0(j, \mathbf{X})]$ as an offset term and weights $w_i(j)$. Denote the estimated intercept term as $\hat{\epsilon}$
 - (4) Let $Q^1(j, \mathbf{X}_i) = \text{expit}\{\text{logit}[Q^0(j, \mathbf{X}_i)] + \hat{\epsilon}\}$

The TMLE estimator is then

$$\hat{\tau}_{t,t'}^{TMLE} = \frac{1}{n} \sum_{i=1}^n [Q^1(t, \mathbf{X}_i) - Q^1(t', \mathbf{X}_i)].$$

If Y is a continuous variable, it is recommended to rescale Y and $Q^0(j, \mathbf{X})$ such that values lie between 0 and 1 before performing step 3 [20,39, Chapter 7]. Let a and b be the known limits of Y , then the rescaled outcome and expectations are $Y^* = \frac{Y-a}{b-a}$ and $Q^{0,*}(j, \mathbf{X}) = \frac{Q^0(j, \mathbf{X})-a}{b-a}$. The causal effect estimate $\hat{\tau}_{t,t'}^{TMLE}$ is computed after back-transforming $Q^1(t, \mathbf{X}_i)$ and $Q^1(t', \mathbf{X}_i)$ into the original scale.

An asymptotic estimator for the variance of $\hat{\tau}_{t,t'}^{TMLE}$ is based on the sample variance of the efficient influence function (see, for example, [39, Chapter 5]). This variance estimator is consistent when both of these models are estimated with consistent estimators and asymptotically conservative when only the treatment model estimator is consistent [39].

2.5. Overlap weights

When using the overlap weights, each observation is attributed a weight $w_i = h(\mathbf{X}_i)/\hat{P}(T = t_i|\mathbf{X}_i)$, where $h(\mathbf{X}_i) = [\sum_{l=1}^k 1/\hat{P}(T = l|\mathbf{X}_i)]^{-1}$. The overlap weight estimator is then

$$\hat{\tau}_{t,t'}^{OW} = \frac{\sum_{i=1}^n I(T_i = t)Y_iw_i}{\sum_{i=1}^n I(T_i = t)w_i} - \frac{\sum_{i=1}^n I(T_i = t')Y_iw_i}{\sum_{i=1}^n I(T_i = t')w_i}.$$

This specific choice of $h(\mathbf{X}_i)$ yields an estimator with minimal variance, assuming homoscedastic residual variances of the potential outcomes [22].

It is important to note that $\hat{\tau}_{t,t'}^{OW}$ is not generally a consistent estimator for the population average pairwise treatment effect $\tau_{t,t'}$. In fact, the estimand of $\hat{\tau}_{t,t'}^{OW}$ is the so-called average pairwise treatment effect in the overlap population:

$$\tau_{t,t'}^* = \frac{\int_{\mathbf{X} \in \mathcal{X}} \{ \mathbb{E}[Y(t)|\mathbf{X}] - \mathbb{E}[Y(t')|\mathbf{X}] \} f(\mathbf{X})h(\mathbf{X})\mu(d\mathbf{X})}{\int_{\mathbf{X} \in \mathcal{X}} f(\mathbf{X})h(\mathbf{X})\mu(d\mathbf{X})},$$

where $f(\mathbf{X})$ is the marginal density of \mathbf{X} with respect to an appropriate measure μ . This corresponds to a weighted average treatment effect, where $h(\mathbf{X})$ is the weight function. The estimand of the overlap weights focuses on the treatment effect in the part of the population where the positivity assumption holds. Indeed, parts of the populations for which at least one of the treatment is impossible have a weight of 0 in the estimand.

A variance estimator for $\hat{\tau}_{t,t'}^{OW}$ that accounts for the estimation of the weights is provided by [22]. Because this variance estimator requires computing derivatives of the weights according to parameters of the treatment model, it combines poorly with machine learning methods, which may lack a simple parametric expression. We thus propose an alternative variance estimator similar to the variance estimator commonly employed for inverse probability weighting estimators. This simpler conservative variance estimator for $\hat{\tau}_{t,t'}^{OW}$ is obtained by considering the weights as known. It corresponds to the empirical variance of the nonparametric influence function of $\tau_{t,t'}^*$ (given by [28]), scaled by $1/n$:

$$\widehat{Var}(\hat{\tau}_{t,t'}^{OW}) = \frac{1}{n^2} \sum_{i=1}^n \left[\frac{I(T_i = t)(Y_i - \hat{\tau}_t)w_i}{\frac{1}{n} \sum_{i=1}^n h(\mathbf{X}_i)} - \frac{I(T_i = t')(Y_i - \hat{\tau}_{t'})w_i}{\frac{1}{n} \sum_{i=1}^n h(\mathbf{X}_i)} \right]^2,$$

where $\hat{\tau}_t = \frac{\sum_{i=1}^n I(T_i=t)Y_iw_i}{\sum_{i=1}^n I(T_i=t)w_i}$ is the treatment t weighted mean. This variance estimator is motivated by the fact the behavior of a semiparametric estimator is asymptotically the same as the one from its influence function [40, Chapter 5]. An important advantage of our proposed variance estimator is that effect estimates and inferences based on this variance estimator correspond to those produced by standard software when running a weighted generalized estimating equation regression with the treatment as the sole predictor and employing the robust variance estimator (for example, using `proc genmod` in SAS with a `repeated` statement, or using `geeglm` in the `geepack` package in R [11]).

As indicated by [22] it is possible to construct an outcome-regression augmented version of the overlap weight estimator:

$$\hat{\tau}_{t,t'}^{A-OW} = \hat{\tau}_{t,t'}^{OW} + \frac{\sum_{i=1}^n h(\mathbf{X}_i) \left[\hat{\mathbb{E}}(Y|T = t, \mathbf{X}_i) - \hat{\mathbb{E}}(Y|T = t', \mathbf{X}_i) \right]}{\sum_{i=1}^n h(\mathbf{X}_i)} - \frac{\sum_{i=1}^n I(T_i = t) \hat{\mathbb{E}}(Y|T = t, \mathbf{X}_i) w_i}{\sum_{i=1}^n I(T_i = t) w_i} + \frac{\sum_{i=1}^n I(T_i = t') \hat{\mathbb{E}}(Y|T = t', \mathbf{X}_i) w_i}{\sum_{i=1}^n I(T_i = t') w_i}.$$

This estimator is semiparametric efficient for estimating $\tau_{t,t'}^*$ when both the treatment model $\hat{P}(T = t|\mathbf{X})$ and the outcome model $\hat{\mathbb{E}}(Y|T, \mathbf{X}_i)$ are correctly specified [22].

Following [28], we can show that $\hat{\tau}_{t,t'}^{A-OW}$ has a property *analogue* to double-robustness.

Theorem 2.1: *The estimator $\hat{\tau}_{t,t'}^{A-OW}$ is consistent for $\tau_{t,t'}^*$ when the treatment model is correctly specified, whether the outcome model is correctly specified or not. When the outcome model is correctly specified, but the treatment model is misspecified, $\hat{\tau}_{t,t'}^{A-OW}$ is consistent for*

$$\tilde{\tau}_{t,t'}^* = \frac{\int_{\mathbf{X} \in \mathcal{X}} \{ \mathbb{E}[Y(t)|\mathbf{X}] - \mathbb{E}[Y(t')|\mathbf{X}] \} f(\mathbf{X}) \tilde{h}(\mathbf{X}) \mu(d\mathbf{X})}{\int_{\mathbf{X} \in \mathcal{X}} f(\mathbf{X}) \tilde{h}(\mathbf{X}) \mu(d\mathbf{X})},$$

where $\tilde{h}(\mathbf{X})$ is the estimand of $[\sum_{l=1}^k 1/\hat{P}(T = l|\mathbf{X}_i)]^{-1}$ under the misspecified $\hat{P}(T = l|\mathbf{X}_i)$.

The proof of Theorem 2.1 is provided in Web Appendix A.

Based on the efficient influence function derived by [28], we propose an asymptotic variance estimator for $\hat{\tau}_{t,t'}^{A-OW}$ obtained by computing the empirical variance of the nonparametric efficient influence function, scaled by a factor $1/n$

$$\widehat{Var}(\hat{\tau}_{t,t'}^{A-OW}) = \frac{1}{n^2} \sum_{i=1}^n \left\{ \frac{h(\mathbf{X}_i)}{\frac{1}{n} \sum_{i=1}^n h(\mathbf{X}_i)} \left[\hat{\mathbb{E}}(Y|T = t, \mathbf{X}_i) - \hat{\mathbb{E}}(Y|T = t', \mathbf{X}_i) + \frac{I(T_i = t) \{ Y_i - \hat{\mathbb{E}}(Y|T = t, \mathbf{X}_i) \}}{\hat{P}(T = t|\mathbf{X}_i)} - \frac{I(T_i = t') \{ Y_i - \hat{\mathbb{E}}(Y|T = t', \mathbf{X}_i) \}}{\hat{P}(T = t'|\mathbf{X}_i)} \right] - \hat{\tau}_{t,t'}^{A-OW} \right\}^2.$$

This variance estimator is consistent if both the treatment and outcome models are estimated with consistent estimators. Like the variance estimator for $\hat{\tau}_{t,t'}^{OW}$ we proposed, this variance estimator can be conveniently computed when machine learning methods are employed for fitting the treatment or outcome model. Indeed, the variance formulas are relatively simple functions of observed variables and predicted values, thus not directly depending on how the model was fitted.

3. A simulation study

We now describe the simulation study we have conducted to compare the empirical performance of the methods presented in the previous section.

3.1. Monte Carlo simulation

3.1.1. Simulation design and scenarios

For our data-generating process, we considered three covariates X_1, X_2, X_3 arbitrarily generated as follow: $X_1 \sim \mathcal{N}(0, 1)$, $X_2 \sim \mathcal{N}(2X_1, 1)$ and $X_3 \sim \text{Bernoulli}(0.4)$. The treatment variable T was simulated according to a multinomial logistic regression with three levels (1, 2 and 3), where the first was the reference category. The probability of membership in each group was given by:

$$\begin{aligned}\pi_1 &= \left[1 + \exp(\beta_{21}X_1 + \beta_{22}X_2 + \beta_{23}X_3 + \beta_{24}X_1X_3 + \beta_{25}X_1^2) \right. \\ &\quad \left. + \exp(\beta_{31}X_1 + \beta_{32}X_2 + \beta_{33}X_3 + \beta_{34}X_1X_3 + \beta_{35}X_1^2) \right]^{-1}, \\ \pi_2 &= \exp(\beta_{21}X_1 + \beta_{22}X_2 + \beta_{23}X_3 + \beta_{24}X_1X_3 + \beta_{25}X_1^2)\pi_1, \\ \pi_3 &= \exp(\beta_{31}X_1 + \beta_{32}X_2 + \beta_{33}X_3 + \beta_{34}X_1X_3 + \beta_{35}X_1^2)\pi_1.\end{aligned}$$

The outcome Y was generated from a normal distribution $Y \sim \mathcal{N}(\mu_3, 1)$, where $\mu_3 = \gamma_0 + \gamma_1X_1 + \gamma_2X_2 + \gamma_3X_3 + \gamma_4X_2X_3 + \gamma_5X_2^2 + \lambda_1I(T = 2) + \lambda_2I(T = 3)$.

We defined four simulation scenarios where the association between the covariates and the treatment was either weak ($T-$) or strong ($T+$), and the association between the covariates and the outcome was either weak ($Y-$) or strong ($Y+$). Parameters λ_1 and λ_2 represent the true treatment effects of level 2 and 3 as compared to level 1 (τ_{21} and τ_{31}), respectively, and were set to $\lambda_1 = 1.0$ and $\lambda_2 = 1.5$. These Monte Carlo scenarios were devised such that the estimand of the overlap weights is the same as the one from the other methods. Such a choice facilitates the comparison between estimators. However, in the plasmode simulation presented in Section 3.2, the estimand of the overlap weight estimators does not correspond with that of the other methods.

The specific choice of the values for β and γ was made to attain the following desiderata: (1) in Scenario $T - Y-$, the confounding bias for both λ parameters is between 10% and 30%, (2) the bias for Scenarios $T - Y+$ and $T + Y-$ is between 40% and 60%, (3) the bias for Scenario $T + Y+$ is above 100%, (4) there is a good overlap of the treatment probabilities' distribution between treatment groups in Scenarios $T-$, and (5) there is a poor overlap in Scenarios $T+$ (near positivity violation).

In Scenarios $T-$, the β values were $\beta_{21} = -0.2, \beta_{22} = 0.2, \beta_{23} = 0.2, \beta_{24} = 0.1, \beta_{25} = 0.1, \beta_{31} = 0.2, \beta_{32} = 0.1, \beta_{33} = -0.2, \beta_{34} = 0.1, \beta_{35} = 0.1$. In Scenarios $T+$, they were $\beta_{21} = -0.8, \beta_{22} = 0.8, \beta_{23} = 0.8, \beta_{24} = 0.2, \beta_{25} = 0.2, \beta_{31} = 0.5, \beta_{32} = 0.5, \beta_{33} = -0.8, \beta_{34} = 0.2, \beta_{35} = 0.2$. In Scenarios $Y-$, the γ values were $\gamma_0 = 0, \gamma_1 = 0.2, \gamma_2 = 0.2, \gamma_3 = 0.2, \gamma_4 = 0.1$ and $\gamma_5 = 0.1$. In Scenarios $Y+$, they were $\gamma_0 = 0, \gamma_1 = 0.5, \gamma_2 = 0.5, \gamma_3 = 0.5, \gamma_4 = 0.2$ and $\gamma_5 = 0.2$.

For each scenario, we simulated 1000 independent data sets each of sample size 500, 1000 and 2000 under the above conditions. We then estimated τ_{21} and τ_{31} utilizing the standardization (*stan*), inverse probability weighting (*IPW*), matching (*match*), bias-corrected matching (*BCM*), targeted maximum likelihood (*TMLE*), overlap weights (*OW*) and augmented overlap weights (*A-OW*) estimators presented previously. For *match* and *BCM*, the matching was performed on the logit of the estimated generalized propensity score (the vector $\text{logit}[\hat{P}(T = t|X)]$). First, the estimators were implemented using correctly specified parametric models, that is, a linear regression of Y on T, X_1, X_2, X_3, X_2X_3 and X_2^2 , and a

multinomial logistic regression of T on X_1, X_2, X_3, X_1X_3 and X_1^2 . This allows investigating the performance of the estimators under ideal circumstances. Next, the estimators were implemented using parametric models that include only main terms, excluding interaction or quadratic terms. This corresponds to a common implementation since the correct model is unknown in practice. Finally, the estimators were implemented using machine learning approaches. For the outcome, we used a Super Learner. The Super Learner is an ensemble method that yields an estimate corresponding to a weighted average of the prediction of multiple procedures, where the weights are determined using cross-validation [38]. The prediction procedures we used were a main terms only linear regression, a main, interactions and quadratic terms linear regression, and a generalized additive model with cubic splines. This was performed using the `SuperLearner` package in R [31]. Because this package does not accommodate multinomial dependent variables, we used a polychotomous regression and multiple classification with the R package `polyspline` [18,19] to model the exposure. This approach uses linear splines and their tensor products to produce predictions. The machine learning methods we chose were selected such that their rate of convergence was \sqrt{n} .

We computed the following measures to assess and compare the performance of the estimators: bias, standard deviation (Std), ratio of the mean estimated standard error to the standard deviation of the estimates, square-root of the mean-squared error (RMSE) and the proportion of the replications in which 95% confidence intervals included the true value of the treatment effect (Coverage CI). For *stan*, inferences were produced using a nonparametric bootstrap estimate of the variance with 200 samples. For *IPW*, *TMLE*, *OW*, and *A-OW*, the influence function variance estimator was used. For *match* and *BCM*, we used the variance estimator proposed by [35]. To assess bias levels, we calculated an unadjusted estimate of the treatment effect.

3.1.2. Simulation results

Tables 1, 2, 3 and 4 compare the bias, standard deviation, RMSE and coverage of confidence intervals across methods for $n = 1000$. Similar results were obtained for $n = 500$ and $n = 2000$. A detailed presentation of these results is available in the Supplemental Online Material. Figure 1 depicts the ratio of the mean estimated standard error to the standard deviation of the estimates of τ_{31} . Similar results were observed for τ_{21} and are available in the Supplemental Online Material. A figure depicting the overlap in the treatment probabilities distribution according to treatment group and scenarios is also available in the Supplemental Online Material.

Under a correct parametric specification, all methods managed to almost completely eliminate the bias in Scenarios $T - Y-$ and $T - Y+$. In Scenarios with positivity violations ($T + Y-$ and $T + Y+$), *IPW* and *match* yielded results with substantial residual bias, while the other methods still achieved close to unbiased estimation. In all scenarios, *stan* produced the estimates with lowest standard deviation and RMSE, followed closely by *A-OW*, and by *TMLE* in Scenarios $T - Y-$ and $T - Y+$. In Scenarios $T + Y-$ and $T + Y+$, the standard deviation and RMSE of *IPW* and *match* were drastically larger than those of the other methods. The coverage of confidence intervals of all methods was close to 95% or slightly conservative, except for *IPW* and *match* in Scenarios $T + Y-$ and $T + Y+$ where confidence intervals included the true effect much less often than expected.

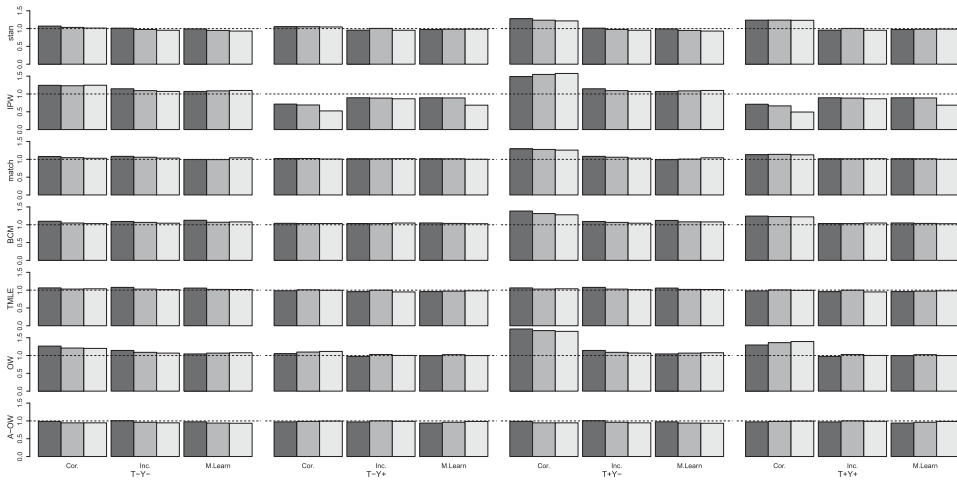


Figure 1. Ratio of the mean estimated standard error to the standard deviation of the estimates of τ_{31} for the correct parametric implementation according to sample size (dark gray = 500, gray = 1000, light gray = 2000), the strength of the association between covariates and treatment (columns; weak= $T -$ or strong= $T +$) and between the covariates and outcome (weak= $Y -$ or strong= $Y +$), and implementation (rows; Cor = Correct parametric, Inc = Incorrect parametric, ML = machine learning).

When using an incorrect parametric specification, important bias reduction was still achieved, but residual confounding bias remained for all adjustment methods. In Scenarios $T - Y -$ and $T - Y +$, BCM and $match$ were the methods that reduced the bias the most and were the only methods that yielded confidence intervals with appropriate coverage. In Scenarios $T + Y -$ and $T + Y +$, $A - OW$ and BCM were the methods that produced the most important bias reduction and the coverage of their confidence intervals were the closest, yet inferior, to 95%.

When a machine learning implementation was used, $stan$, BCM , $TMLE$ and $A - OW$ produced results very similar to those obtained under the correct parametric implementations, with regards to bias, standard deviation, RMSE and coverage of confidence intervals. However, IPW , $match$ and OW produced biased results with increased standard deviations and RMSE. The coverage of their confidence intervals was also generally below the expected level, except for $match$ that retained appropriate coverage in Scenarios $T - Y -$ and $T - Y +$.

In all scenarios, we noticed that the bias of $match$ and BCM tended to decrease with increased sample size, even under an incorrect parametric implementation.

Figure 1 indicates that the variance estimators we proposed for OW and $A - OW$ performed well. As expected, the variance estimator for OW tends to overestimate the true variance. The estimated standard error of $A - OW$ was very similar to its Monte Carlo standard deviation. Figure 1 also reveals that the usual sandwich estimator of the variance of IPW , which is expected to be conservative, sometimes drastically underestimated the true variance in Scenarios $T - Y +$ and $T + Y +$. The variance estimators of $stan$, $match$ and BCM were conservative in Scenarios $T + Y -$ and $T + Y +$ when using a correct parametric implementation. In all other cases, the estimated and empirical variances were similar.

Table 1. Estimate of the treatment effect in the Scenario with a weak association between covariates and treatment, a weak association between covariates and outcome ($T - Y-$), and a sample size of 1000.

Implementation	Approach	Bias		Std		RMSE		Coverage CI	
		$\hat{\tau}_{21}$	$\hat{\tau}_{31}$	$\hat{\tau}_{21}$	$\hat{\tau}_{31}$	$\hat{\tau}_{21}$	$\hat{\tau}_{31}$	$\hat{\tau}_{21}$	$\hat{\tau}_{31}$
Cor.param.	Crude	0.240	0.347	0.106	0.112	0.262	0.365	0.422	0.129
	stan	0.001	0.002	0.081	0.082	0.081	0.082	0.948	0.960
	IPW	0.003	0.006	0.102	0.103	0.102	0.103	0.983	0.985
	match	0.018	0.019	0.099	0.099	0.101	0.100	0.956	0.961
	BCM	0.003	0.003	0.098	0.098	0.098	0.098	0.958	0.962
	TMLE	0.002	0.003	0.084	0.085	0.084	0.085	0.953	0.960
	OW	-0.000	0.002	0.091	0.092	0.091	0.092	0.985	0.978
	A-OW	0.001	0.002	0.083	0.084	0.083	0.084	0.940	0.943
Inc.param.	stan	0.046	0.057	0.096	0.097	0.107	0.112	0.922	0.905
	IPW	0.092	0.087	0.104	0.104	0.140	0.136	0.882	0.898
	match	0.032	0.034	0.097	0.099	0.102	0.105	0.953	0.955
	BCM	0.024	0.022	0.096	0.098	0.099	0.100	0.959	0.958
	TMLE	0.076	0.077	0.100	0.102	0.126	0.128	0.890	0.902
	OW	0.083	0.078	0.100	0.100	0.130	0.127	0.893	0.908
	A-OW	0.068	0.069	0.097	0.098	0.118	0.120	0.870	0.885
	stan	0.001	0.005	0.084	0.085	0.084	0.085	0.926	0.941
M.Learning	IPW	0.101	0.079	0.106	0.105	0.146	0.132	0.857	0.905
	match	0.035	0.020	0.109	0.110	0.115	0.111	0.940	0.954
	BCM	-0.003	0.001	0.101	0.102	0.101	0.102	0.958	0.965
	TMLE	0.004	0.007	0.085	0.086	0.085	0.086	0.957	0.955
	OW	0.091	0.070	0.102	0.102	0.137	0.124	0.868	0.907
	A-OW	0.002	0.006	0.084	0.085	0.085	0.086	0.933	0.941

Note: Cor.param=correct parametric models, Inc.param=incorrect parametric models, M.Learning=machine learning, Crude=Unadjusted, stan=standardization, IPW=inverse probability weighting, match=matching, BCM=bias-corrected matching, TMLE=targeted maximum likelihood, OW=overlap weights, A-OW=augmented overlap weights.

3.2. Plasmode simulation

3.2.1. Plasmode datasets

We used the data on forest fires in Alberta, Canada, as a basis for our plasmode simulation. These data are produced and published online by the Government of Alberta [41]. The purpose of the analysis was to compare various interventions for fighting wildfires in Alberta on their probability of preventing the fire to grow after its initial assessment. We considered only fires caused by lightning from 1996 to 2014. Observations for which size at ‘being held’ was smaller than the size at initial attack were also removed. ‘Being held’ is defined as in [37] as a state when no further increase in size is expected.

Potential confounders considered in the plasmode simulation are the ecological region in which the fire occurred (Clear Hills Upland, Mid-Boreal Uplands, Other), the number of fires active at the time of initial assessment of each fire, fuel type (Boreal Spruce, Boreal Mixedwood – Green, other), month of the year the fire was first assessed (‘May or June’, July, ‘August, September or October’), how the fire was discovered (air patrol, lookout, unplanned), response time in hours between the moment the fire was reported to first assessment by forestry personnel, and the natural logarithm (\ln) of the size of the fire at the time of the initial attack. Response time was truncated at the 95th percentile of the distribution to limit the influence of extreme observations, which caused convergence problems of models in some replications of the simulation. Additional potential confounders were available but were not considered in the plasmode simulation. As will be explained shortly,

Table 2. Estimate of the treatment effect in the Scenario with a strong association between covariates and treatment, a weak association between covariates and outcome ($T + Y-$), and a sample size of 1000.

Implementation	Approach	Bias		Std		RMSE		Coverage IC		
		$\hat{\tau}_{21}$	$\hat{\tau}_{31}$	$\hat{\tau}_{21}$	$\hat{\tau}_{31}$	$\hat{\tau}_{21}$	$\hat{\tau}_{31}$	$\hat{\tau}_{21}$	$\hat{\tau}_{31}$	
Cor.param.	Crude	0.467	0.790	0.094	0.117	0.477	0.799	0.001	0.000	
	stan	0.005	0.001	0.083	0.094	0.083	0.094	0.955	0.967	
	IPW	0.039	0.032	0.301	0.308	0.303	0.310	0.870	0.898	
	match	0.102	0.092	0.121	0.140	0.158	0.168	0.862	0.889	
	BCM	0.022	0.004	0.120	0.137	0.122	0.137	0.958	0.960	
	TMLE	0.006	0.002	0.101	0.118	0.102	0.118	0.952	0.960	
	OW	0.009	0.004	0.104	0.116	0.105	0.116	0.967	0.975	
	A-OW	0.006	0.003	0.096	0.107	0.096	0.107	0.948	0.959	
	Inc.param.	stan	-0.121	-0.039	0.098	0.113	0.155	0.120	0.785	0.928
		IPW	0.209	0.181	0.174	0.195	0.272	0.266	0.621	0.755
match		0.132	0.118	0.121	0.136	0.179	0.180	0.799	0.869	
BCM		0.091	0.045	0.119	0.136	0.149	0.143	0.887	0.940	
TMLE		0.159	0.112	0.131	0.155	0.206	0.191	0.767	0.892	
OW		0.117	0.099	0.109	0.118	0.160	0.154	0.812	0.865	
A-OW		0.065	0.062	0.101	0.116	0.120	0.131	0.898	0.920	
stan		0.005	0.014	0.090	0.105	0.090	0.106	0.938	0.940	
M.Learning	IPW	0.159	0.138	0.185	0.200	0.244	0.244	0.720	0.801	
	match	0.124	0.111	0.122	0.138	0.174	0.177	0.819	0.875	
	BCM	0.010	0.008	0.118	0.135	0.119	0.136	0.953	0.956	
	TMLE	0.016	0.006	0.106	0.123	0.108	0.124	0.939	0.945	
	OW	0.080	0.067	0.110	0.121	0.136	0.138	0.893	0.912	
	A-OW*	0.003	-0.001	0.098	0.108	0.098	0.108	0.937	0.952	

Note: Cor.param=correct parametric models, Inc.param=incorrect parametric models, M.Learning=machine learning, Crude=Unadjusted, stan=standardization, IPW=inverse probability weighting, match=matching, BCM=bias-corrected matching, TMLE=targeted maximum likelihood, OW=overlap weights, A-OW=augmented overlap weights. **: in 1 replication, the confidence intervals could not be computed.

the plasmode simulation was run on a reduced sample of the total data available. As such, including all potential confounders in the simulations caused convergence issues.

Some levels of categorical variables having few instances were dropped, and observations in those categories were removed. The final database contained 8591 observations, the 7 covariates presented above, and the treatment variable indicating the method of intervention used by the firefighters to suppress the wildfire and the outcome variable. The categories for the treatment were heli-attack crew with helicopter but no rappel capability (HAC1H; 53.8%), heli-attack crew with helicopter and rappel capability (HAC1R; 15.8%), fire-attack crew with or without a helicopter and no rappel capability (HAC1F; 6.3%), Air tanker (15.3%) and Ground-based action (8.8%). The binary response variable was 1 or 0 depending on whether the fire did or did not increase in size between initial attack and 'being held' ($n=1982$ and 6645 , respectively).

For generating plasmode data, we first fitted a random forest regression of the outcome variable conditional on treatment and all seven covariates using the package `randomForest` in R with the default settings [24]. Similarly, a random forest regression of the treatment according to the covariates was fitted. These models were then used to generate simulated treatment and outcome variables. This approach naturally allows for treatment-covariate interactions in generating the outcome and thus allows the estimand of the overlap weight estimators to differ from that of the other estimators. The true causal effects were estimated using a Monte Carlo simulation. Briefly, for each observation in the complete population, we generated counterfactual outcomes that would have been observed under each possible treatment according to the true outcome model. The

Table 3. Estimate of the treatment effect in the Scenario with a weak association between covariates and treatment, a strong association between covariates and outcome ($T - Y_+$), and a sample size of 1000.

Implementation	Approach	Bias		Std		RMSE		Coverage IC	
		$\hat{\tau}_{21}$	$\hat{\tau}_{31}$	$\hat{\tau}_{21}$	$\hat{\tau}_{31}$	$\hat{\tau}_{21}$	$\hat{\tau}_{31}$	$\hat{\tau}_{21}$	$\hat{\tau}_{31}$
Cor.param.	Crude	0.564	0.819	0.173	0.192	0.590	0.841	0.125	0.009
	stan	0.001	0.002	0.081	0.082	0.081	0.082	0.984	0.986
	IPW	0.005	0.008	0.149	0.151	0.149	0.151	0.997	0.995
	match	0.037	0.039	0.110	0.110	0.116	0.117	0.983	0.979
	BCM	0.004	0.002	0.106	0.107	0.106	0.107	0.985	0.990
	TMLE	0.002	0.003	0.084	0.085	0.084	0.085	0.953	0.960
	OW	-0.001	0.002	0.111	0.115	0.111	0.115	0.999	0.998
	A-OW	0.001	0.002	0.083	0.084	0.083	0.084	0.940	0.943
Inc.param.	stan	0.046	0.057	0.096	0.097	0.107	0.112	0.922	0.905
	IPW	0.092	0.087	0.104	0.104	0.140	0.136	0.882	0.898
	match	0.032	0.034	0.097	0.099	0.102	0.105	0.953	0.955
	BCM	0.024	0.022	0.096	0.098	0.099	0.100	0.959	0.958
	TMLE	0.076	0.077	0.100	0.102	0.126	0.128	0.890	0.902
	OW	0.083	0.078	0.100	0.100	0.130	0.127	0.893	0.908
	A-OW	0.068	0.069	0.097	0.098	0.118	0.120	0.870	0.885
	stan	0.001	0.005	0.084	0.085	0.084	0.085	0.926	0.941
M.Learning	IPW	0.101	0.079	0.106	0.105	0.146	0.132	0.857	0.905
	match	0.034	0.019	0.108	0.108	0.113	0.110	0.940	0.954
	BCM	-0.003	0.001	0.099	0.101	0.099	0.101	0.960	0.965
	TMLE	0.004	0.007	0.085	0.086	0.085	0.086	0.957	0.955
	OW	0.091	0.070	0.102	0.102	0.137	0.124	0.868	0.907
	A-OW	0.002	0.006	0.084	0.085	0.085	0.086	0.933	0.941

Note: Cor.param=correct parametric models, Inc.param=incorrect parametric models, M.Learning=machine learning, Crude=Unadjusted, stan=standardization, IPW=inverse probability weighting, match=matching, BCM=bias-corrected matching, TMLE=targeted maximum likelihood, OW=overlap weights, A-OW=augmented overlap weights.

difference in risk of fire progression were then computed for each treatment level, as compared to HAC1H, which was used as the reference level. For the overlap weights, risk differences weighted according to $h(\mathbf{X})$ as defined in Section 2.5 were computed. The true effects comparing Air tanker, Ground-based action, HAC1F and HAC1R to HAC1H in the entire population were, respectively, 0.099, -0.011, 0.012 and -0.003. In the overlap population, the true risk differences were 0.101, -0.024, 0.003 and -0.002. To simplify the presentation, both sets of parameters will be denoted as τ_{21} , τ_{31} , τ_{41} and τ_{51} in the results below.

A plasmode simulated dataset was obtained as follows. First, we drew with replacement a random sample of size 2000 from the original data. Then, for each observation, a new simulated treatment and a new simulated outcome were generated, with $P(T = t|\mathbf{X})$ and $P(Y = 1|T, \mathbf{X})$ determined by the fitted random forest models. Based on this process, we simulated 1000 independent data sets and compared adjustments methods using the same performance metrics as in the Monte Carlo simulation (see Section 3.1.1). For each method, we considered either a parametric models implementation that included only main terms, or a machine learning implementation. For the treatment, these implementations were the same as described in Section 3.1.1. The implementations for the outcome were also similar to those described in Section 3.1.1, but replacing linear regressions by logistic regressions. Since the outcome and treatment data were generated nonparametrically, the correct models are unknown.

Table 4. Estimate of the treatment effect in the Scenario with a strong association between covariates and treatment, a strong association between covariates and outcome ($T + Y+$), and a sample size of 1000.

Implementation	Approach	Bias		Std		RMSE		Coverage IC	
		$\hat{\tau}_{21}$	$\hat{\tau}_{31}$	$\hat{\tau}_{21}$	$\hat{\tau}_{31}$	$\hat{\tau}_{21}$	$\hat{\tau}_{31}$	$\hat{\tau}_{21}$	$\hat{\tau}_{31}$
Cor.param.	Crude	1.180	1.937	0.146	0.198	1.189	1.947	0.000	0.000
	stan	0.005	0.001	0.083	0.094	0.083	0.094	0.987	0.988
	IPW	0.079	0.075	0.612	0.614	0.617	0.618	0.841	0.866
	match	0.213	0.214	0.143	0.161	0.256	0.267	0.747	0.796
	BCM	0.039	0.007	0.131	0.147	0.136	0.147	0.976	0.983
	TMLE	0.006	0.002	0.102	0.118	0.102	0.118	0.951	0.959
	OW	0.012	0.006	0.131	0.142	0.131	0.143	0.990	0.990
	A-OW	0.006	0.003	0.096	0.107	0.096	0.107	0.948	0.959
	Inc.param.	stan	-0.121	-0.039	0.098	0.113	0.155	0.120	0.785
IPW	0.209	0.181	0.174	0.195	0.272	0.266	0.621	0.755	
match	0.132	0.118	0.121	0.136	0.179	0.180	0.799	0.869	
BCM	0.091	0.045	0.119	0.136	0.149	0.143	0.887	0.940	
TMLE	0.159	0.112	0.131	0.155	0.206	0.191	0.767	0.892	
OW	0.117	0.099	0.109	0.118	0.160	0.154	0.812	0.865	
A-OW	0.065	0.062	0.101	0.116	0.120	0.131	0.898	0.920	
M.Learning	stan	0.005	0.014	0.090	0.105	0.090	0.106	0.938	0.940
	IPW	0.159	0.138	0.185	0.200	0.244	0.244	0.720	0.801
	match	0.124	0.111	0.122	0.138	0.174	0.177	0.819	0.875
	BCM	0.010	0.008	0.118	0.135	0.119	0.136	0.953	0.956
	TMLE	0.016	0.006	0.106	0.123	0.108	0.124	0.939	0.945
	OW	0.080	0.067	0.110	0.121	0.136	0.138	0.893	0.912
	A-OW*	0.003	-0.001	0.098	0.108	0.098	0.108	0.937	0.952

Note: Cor.param=correct parametric models, Inc.param=incorrect parametric models, M.Learning=machine learning, Crude=Unadjusted, stan=standardization, IPW=inverse probability weighting, match=matching, BCM=bias-corrected matching, TMLE=targeted maximum likelihood, OW=overlap weights, A-OW=augmented overlap weights. *: in 1 replication, the confidence intervals could not be computed.

3.2.2. Simulation results

A figure displaying the distribution of the treatment probabilities according to treatment groups is available in the Supplemental Online Material. While a good level of overlap between treatment groups was present, multiple observations had treatment probabilities close to 0, thus suggesting possible practical positivity violations.

The results presented in Table 5 show that a large amount of bias is present when no adjustment is performed, especially for τ_{21} . When using a parametric implementation, most methods achieved a bias reduction for estimating all parameters, except *stan* and *match* that increased the bias for one parameter. Additionally, a substantial bias remained for estimating some parameters for *stan*, *match*, and *A-OW*. The RMSE of *stan*, *TMLE*, and *OW* were smaller than those of the crude estimates for all parameters, whereas *IPW*, *match*, *BCM*, and *A-OW* had an increased RMSE for at least one parameter. The lowest RMSE for all parameters was produced by *stan*. Most methods yielded 95% confidence intervals with coverage rate close to the expected level for τ_{21} , τ_{41} and τ_{51} , but not for τ_{31} . Only *IPW* and *OW* had close to adequate coverage for τ_{31} . *match* and *BCM* yielded inadequate confidence intervals for all parameters, because of an underestimation of the true variance (see Figure 3 in the Supplemental Online Material). When using a machine learning implementation, the performance of most adjustment methods was overall marginally improved. Indeed, the bias and RMSE were generally smaller, and coverage closer to 95%. For *IPW*, however, the results were somewhat worse.

Table 5. Estimate of treatment effect in plasmode simulation using 2000 observations. True effects are $\tau_{21} = -0.099$, $\tau_{31} = 0.011$, $\tau_{41} = -0.012$ and $\tau_{51} = 0.003$ for all methods except *OW* and *A-OW* and $\tau_{21} = -0.101$, $\tau_{31} = 0.024$, $\tau_{41} = -0.003$ and $\tau_{51} = 0.002$ for the these two methods. HAC1H is the reference category . τ_{21} , τ_{31} , τ_{41} and τ_{51} refer to treatment effect associated to Air tanker, Ground-based action, HAC1F and HAC1R, respectively.

Implementation	Approach	Bias				RMSE				Coverage IC			
		$\hat{\tau}_{21}$	$\hat{\tau}_{31}$	$\hat{\tau}_{41}$	$\hat{\tau}_{51}$	$\hat{\tau}_{21}$	$\hat{\tau}_{31}$	$\hat{\tau}_{41}$	$\hat{\tau}_{51}$	$\hat{\tau}_{21}$	$\hat{\tau}_{31}$	$\hat{\tau}_{41}$	$\hat{\tau}_{51}$
Parametric	Crude	-0.164	-0.027	-0.021	-0.035	0.168	0.045	0.050	0.045	0.002	0.897	0.941	0.786
	stan	0.004	0.028	0.009	-0.005	0.029	0.043	0.041	0.026	0.921	0.800	0.939	0.942
	IPW	-0.005	0.011	0.003	-0.002	0.040	0.048	0.050	0.029	0.957	0.902	0.947	0.963
	match	0.040	-0.004	0.023	0.004	0.050	0.043	0.052	0.042	0.279	0.428	0.335	0.398
	BCM	0.004	0.003	0.002	-0.013	0.045	0.052	0.053	0.055	0.406	0.353	0.324	0.326
	TMLE	-0.006	0.012	0.001	-0.005	0.037	0.044	0.044	0.028	0.935	0.859	0.926	0.951
	OW	0.001	0.005	0.009	0.004	0.041	0.045	0.049	0.033	0.955	0.902	0.936	0.966
	A-OW	-0.006	0.021	0.007	-0.005	0.040	0.048	0.046	0.032	0.942	0.859	0.943	0.942
Machine learning	stan	0.007	0.027	0.008	-0.006	0.030	0.043	0.041	0.026	0.923	0.847	0.941	0.944
	IPW	-0.027	0.009	-0.001	-0.012	0.048	0.040	0.046	0.031	0.908	0.924	0.956	0.954
	match	0.028	-0.008	0.019	0.000	0.041	0.037	0.048	0.039	0.401	0.485	0.360	0.459
	BCM	0.001	0.007	0.006	-0.007	0.042	0.044	0.048	0.050	0.436	0.393	0.349	0.367
	TMLE	-0.003	0.018	0.007	-0.004	0.035	0.043	0.042	0.027	0.932	0.837	0.925	0.942
	OW	-0.028	0.000	-0.007	-0.011	0.047	0.040	0.048	0.032	0.899	0.916	0.952	0.959
	A-OW	-0.005	0.025	0.009	-0.004	0.035	0.047	0.045	0.028	0.935	0.824	0.930	0.948

Note: Crude=Unadjusted, stan=standardization, IPW=inverse probability weighting, match=matching, BCM=bias-corrected matching, TMLE=targeted maximum likelihood, OW=overlap weights, A-OW=augmented overlap weights.

Figure 3 in the Supplemental Online Material represents the ratio of the mean estimated standard error to the standard deviation of the estimates. The variances estimated using our proposed estimators for *OW* and *A-OW* were similar to the Monte Carlo variances. In fact, all variance estimators except those of *match* and *BCM* performed well in the plasmode simulation.

4. Application

4.1. Context

Wildfires have great economic, environmental and societal impacts. A key tool of fire management is fire suppression by initial attack teams who seek to limit the growth, and hence the final size, of the fire [7]. These interventions by initial attack differ in terms of the size and training of the fire-fighting crews, and the mechanical resources provided to them. However, few studies have attempted estimating the causal effect of initial attack on wildfire growth.

Recently, [37] compared initial attack interventions using public data on wildfire in Alberta, Canada. More aggressive interventions were associated with greater fire growth, opposite to the expected direction of the causal effect. These results suggest that important confounding by indication may be present in fire management agency records.

4.2. Data

The data we considered are similar to those used for performing the plasmode simulation and those employed by [37]. More precisely, we considered all fires on provincial land or

in other public lands caused by lighting between 2003 and 2014 recorded in the historical wildfire database of Alberta's Agriculture and Forestry ministry [41].

The final data consisted of 5439 observations. The fire-attack interventions being compared are the same as those in the plasmode simulation: heli-attack crew with helicopter but no rappel capability (HAC1H), heli-attack crew with helicopter and rappel capability (HAC1R), fire-attack crew with or without a helicopter and no rappel capability (HAC1F), Air tanker, and Ground-based action. The outcome of interest was whether the fire grew in size between initial assessment and 'being held'.

Based on experts' knowledge and the scientific literature, the following 11 potential confounders were considered: (1) Initial Spread Index, (2) Fire Weather Index, (5) ecological region in which the fire occurred (Clear Hills Upland, Mid-Boreal Uplands, Other), (6) fuel type at initial assessment (Boreal Spruce, Boreal Mixedwood – Green, Other), (7) period of day (AM or PM), (8) month of the year the fire was first assessed ('May or June', July, 'August, September or October'), (9) response time in hours between the moment the fire was reported to first assessment by fire fighters, (10) the number of fires active at the time of initial assessment of each fire, and (11) the natural logarithm (\ln) of the size of the fire at the initial attack. The Initial Spread Index and the Fire Weather Index are used to predict the rate of progression of fires and fire danger, respectively, by Alberta's Wildfire Management Branch. These indexes are notably based on daily weather variables. To account for the fact that initial attack decisions would be based not only on current and recent weather, but also on future forecast, we used the values of these variables for the day of fire assessment, the two days prior to assessment and the two days following assessment (for a total of five variables for each index). More information on these variables is available in [37] and references therein.

4.3. Analysis

Not all interventions would be reasonable choices to suppress some fires given their characteristics at the moment of assessment. As such, it seems more relevant to target the average effect of interventions only for those fires for which multiple options are reasonable. Based on this and on the results from our simulation study, we decided to use the augmented overlap-weights estimator with a machine learning implementation. The treatment probabilities were estimated using a polychotomous regression and the outcome expectations were estimated with the Super Learner using a main term generalized linear model, a generalized additive model with spline terms, and random forests as prediction procedures. Note that we no longer limit the machine learning methods we consider to those having a convergence rate of \sqrt{n} , since the augmented overlap weight estimator we use combines outcome and treatment modeling, thus allowing the use of more flexible machine learning methods. The generalized linear model with interaction and quadratic terms was not considered as a prediction procedure because of the small sample size in some levels of categorical variables. All potential confounders were included as independent variables in both models. Year was considered as a categorical variable (12 levels). Inferences were produced using the variance estimator we introduced in Section 2.5. *P*-values were adjusted for multiple comparisons with Holm's method [13].

Table 6. Association between initial intervention used to suppress the fire and probability of fire growth between initial assessment and being held.

Contrast	Risk difference	95% confidence interval	Adjusted <i>p</i> -value
Air tanker vs HAC1H	0.107	(0.083, 0.131)	< 0.001
Ground-based action vs HAC1H	-0.031	(-0.056, -0.005)	0.121
HAC1F vs HAC1H	-0.008	(-0.031, 0.015)	1.000
HAC1R vs HAC1H	-0.005	(-0.032, 0.022)	1.000
Ground-based action vs Air Tanker	-0.138	(-0.170, -0.105)	< 0.001
HAC1F vs Air Tanker	-0.115	(-0.145, -0.085)	< 0.001
HAC1R vs Air Tanker	-0.112	(-0.145, -0.079)	< 0.001
HAC1F vs Ground-based action	0.023	(-0.009, 0.055)	0.749
HAC1R vs Ground-based action	0.026	(-0.009, 0.061)	0.749
HAC1R vs HAC1F	0.003	(-0.030, 0.036)	1.000

Notes: All estimates are adjusted using the augmented overlap weights estimator for Initial Spread Index, Fire Weather Index, year, how the fire was discovered, ecological region, fuel type, period of day, month of the year, response time, number of fires active, and \ln of the size of the fire at the time of the initial attack. Abbreviations: HAC1H = heli-attack crew with helicopter but no rappel capability, HAC1R = heli-attack crew with helicopter and rappel capability, HAC1F = fire-attack crew with or without a helicopter and no rappel capability. *P*-values are adjusted for multiple comparisons using the Holm method.

4.4. Results

Characteristics of the considered fires according to the initial intervention that was chosen for suppressing the fire are presented in Web Appendix C. Important imbalances between treatment groups were observed for calendar year, how the fire was discovered, ecological region, fuel type, response time, number of active fires and initial size of the fire, and some imbalance was observed for the fire weather index on the day of initial assessment. To the best of our knowledge, no method has been developed to assess post-adjustment balance for the augmented overlap weight estimator.

Adjusted associations between initial interventions for suppressing fires and the probability of fire growth are reported in Table 6. Air tanker, which is the most aggressive intervention, is associated with the largest probability of fire growth. All interventions are associated with probabilities of fire growth similar to or greater than that of ground-based action, which we consider to be the least aggressive intervention. The various heli-attack crew interventions are all associated with similar probabilities of fire growth.

5. Conclusion and discussion

The initial motivation for this work was to compare the effect of five initial attack interventions on wildfire growth. Multiple analytical challenges were expected in attempting to estimate this effect, including important confounding by indication, possible non-positivity and unknown model specifications. We reviewed studies evaluating the empirical performance of adjustment methods and identified those that appeared best suited to address these challenges: the overlap weights, the bias-corrected matching and the targeted maximum likelihood estimation.

Regarding the overlap weight estimator, we demonstrated that when it is augmented with an outcome regression, it benefits from a property analogue to double-robustness in the multi-level treatment case, following the proof of [28] for the binary case. Additionally, we have proposed a simple asymptotic variance estimator for the overlap weights and the

augmented overlap weights based on the semiparametric theory. These variance estimators offer multiple benefits. First, they are expressed as the sample variance of a relatively simple quantity that do not require advanced coding or mathematical skills to compute. This is particularly true for the overlap weights, since estimation and inferences correspond to those produced by common generalized estimating equations routines with a robust variance estimator. A further advantage of our proposed variance estimator is that it is agnostic to the model used for computing the point estimates. Hence, it can be computed as easily whether the point estimates are obtained with parametric regression models or machine learning algorithms.

We have also conducted a simulation study comparing the overlap weights, the augmented overlap weights, the bias-corrected matching and the TMLE estimator, as well as standardization, regular inverse probability weighting and matching as benchmark comparators. Most of these methods had never been compared to one another in a multi-level treatment setting. Our proposed variance estimator for the overlap weight estimators performed well, even with small sample sizes. We also observed that methods that combine outcome and treatment modeling (bias-corrected matching, TMLE and augmented overlap weights) performed better than those solely based on the treatment model (inverse probability weighting, matching and overlap weights) in terms of bias reduction, standard deviation of estimates and RMSE. Under positivity violations, inverse probability weighting and matching were biased and had increased variability, while all other methods were relatively unaffected. We also observed that the variance estimator for the matching and bias-corrected matching estimators of [35] performed poorly in the plasmode simulation. We hypothesize that this is because of the binary nature of the outcome.

Unsurprisingly, when an incorrect parametric implementation was used, all adjustment methods were biased. However, the bias for matching and bias-corrected matching tended to get smaller as sample size increased, whereas the bias remained constant for the other methods. A plausible explanation for this phenomenon is that these matching methods are less model dependent, since they involve imputing the missing counterfactual outcomes by the observed outcomes from subjects in the other treatment groups. As sample size increases, the pool of control subjects get larger, thus allowing observations to be matched with others that are more similar to them. We did not observe that methods that combine outcome and exposure modeling performed better than the others under models misspecification. While it is expected that double-robustness theoretically helps protecting against the bias attributable to model misspecifications, this property requires that at least one of the two models involved is correct for unbiased estimation. When both models are incorrect to some extent, as should arguably be expected in practice, double-robust methods do not necessarily produce estimates with less bias than others [16,20]. However, a theoretical advantage of combining outcome and treatment modeling is the possibility to use machine learning methods that have a relatively slow convergence rate without sacrificing the convergence rate of the treatment effect estimator [17].

The results of our simulation study finally indicate that machine learning methods can be helpful for preventing misspecification bias for some, but not all, adjustment methods. Indeed, for all methods that include an outcome modeling component (standardization, bias-corrected matching, TMLE and augmented overlap weight), the results under a machine learning implementation were very similar to those obtained under the ideal case of a correct parametric specification for all considered performance metrics. As such,

the routine use of machine learning algorithms for performing confounder adjustment may have benefits to help preventing bias without paying a price in terms of power. On the other hand, methods that focus only on modeling the treatment (inverse probability weighting, matching and overlap weight) had considerable bias when employing a machine learning implementation. A possible explanation for these results is that machine learning algorithms may tend to more often predict treatment probabilities close to 0 or 1 because they are better able to fit the observed data, thus exacerbating practical positivity violations.

Standardization was the method that overall performed best, closely followed by the augmented overlap weight in most scenarios and sometimes also by TMLE. In practice, we warn that the choice of an approach should not only be guided by its performance, but also by the question of interest. For example, if the goal is to inform about a decision that would apply at the whole population level, the average effect in the entire population would be the one that is of most interest, whereas if the question is to help the decision in cases where several treatment options are possible, the methods that estimate an effect in the subgroups where there is overlap would likely be more appropriate.

Based on the context of the wildfire growth problem and the results of the simulation study, we decided to estimate the effect of initial attack interventions using the augmented overlap weight estimator, implemented with machine learning algorithms. The adjusted associations we observed were counter-intuitive: the more aggressive the intervention was, the larger was the probability that fires grew. These associations are very unlikely to represent the true causal effect. Based on our simulation study's results, we do not believe these results can be explained by residual confounding due to measured confounders or to non-positivity. We thus hypothesize that residual confounding due to important non-measured confounders is present, despite the fact we have made important efforts to include multiple potential confounders in our analysis. Having accounted for the factors previously shown to affect the probabilities of fire growth after initial attack in this system [3], we are unable to advance a more precise explanation. Another possible explanation for our counterintuitive results may be the presence of large effect modification, that is, the treatment effect may vary widely according to the initial characteristics of the fire. It may thus be the case that the less aggressive interventions are indeed more effective than more aggressive interventions in most situations, but less effective in some specific situations. Dynamic treatment regime methods may thus be interesting to explore in future studies to garner a better understanding of the effect of fire suppression methods[29,30,34].

Ultimately, the findings from this study originate from simulation studies and are thus limited to the contexts that were investigated. However, the plasmode simulation based on real data helped understanding the performance of these approaches in a realistic setting. Bias elimination relies on the fact that all confounders are measured, which arguably never occurs in practice. As such, at least some amount of residual confounding due to unmeasured confounders is inevitably present in real data analyses. This is one plausible explanation for the counter-intuitive associations we have observed in the wildfire analysis. Comparing the effect of fire-fighting interventions using observational data is fraught with multiple challenges. Our work contribute to facing these challenges by helping ruling out residual confounding due to known confounders and positivity violations as possible explanations to the unexpected associations that are observed. We hope our study will stimulate others to attempt facing the remaining challenges to be able to better inform fire managers regarding the best course of action when wildfires occur.

Acknowledgments

This research was funded by the Natural Sciences and Engineering Research Council of Canada grant # 2016-06295. S. A. Diop was also supported by a scholarship of Excellence from the Faculté de sciences et de génie of Université Laval. D Talbot is a Fonds de recherche du Québec - Santé Chercheur-Boursier.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

This work was supported by Fonds de Recherche du Québec – Santé [grant number 265385] and Natural Sciences and Engineering Research Council of Canada [grant number 2016-06295].

ORCID

Denis Talbot  <http://orcid.org/0000-0003-0431-3314>

References

- [1] A. Abadie and G.W. Imbens, *Large sample properties of matching estimators for average treatment effects*, *Econometrica* 74 (2006), pp. 235–267.
- [2] A. Abadie and G.W. Imbens, *Bias-corrected matching estimators for average treatment effects*, *J. Bus. Econ. Stat.* 29 (2011), pp. 1–11.
- [3] M.C. Arienti, S.G. Cumming and S. Boutin, *Empirical models of forest fire initial attack success probabilities: the effects of fuels, anthropogenic linear features, fire weather, and management*, *Canadian J. Forest Res.* 36 (2006), pp. 3155–3166.
- [4] P.C. Austin, *The performance of different propensity-score methods for estimating relative risks*, *J. Clin. Epidemiol.* 61 (2008), pp. 537–545.
- [5] P.C. Austin, *The performance of different propensity-score methods for estimating differences in proportions (risk differences or absolute risk reductions) in observational studies*, *Stat. Med.* 29 (2010), pp. 2137–2148.
- [6] P.C. Austin, P. Grootendorst and G.M. Anderson, *A comparison of the ability of different propensity score models to balance measured variables between treated and untreated subjects: a Monte Carlo study*, *Stat. Med.* 26 (2007), pp. 734–753.
- [7] S. Cumming, *Effective fire suppression in boreal forests*, *Canadian J. Forest Res.* 35 (2005), pp. 772–786.
- [8] P. Feng, X.H. Zhou, Q.M. Zou, M.Y. Fan and X.S. Li, *Generalized propensity score for estimating the average treatment effect of multiple treatments*, *Stat. Med.* 31 (2012), pp. 681–697.
- [9] M. Frölich, *Programme evaluation with multiple treatments*, *J. Econ. Surv.* 18 (2004), pp. 181–224.
- [10] G. L. Gadbury, Q. Xiang, L. Yang, S. Barnes, G. P. Page, D. B. Allison and G. Gibson, *Evaluating statistical methods using plasmode data sets in the age of massive public databases: an illustration using false discovery rates*, *PLoS Genet.* 4 (2008), p. e1000098.
- [11] U. Halekoh, S. Hojsgaard and J. Yan, *The r package geeppack for generalized estimating equations*, *J. Stat. Softw.* 15 (2006), pp. 1–11.
- [12] M. Hernán and J. Robins, *Causal Inference: What If*, Chapman & Hall/CRC, Boca Raton, FL, 2020.
- [13] S. Holm, *A simple sequentially rejective multiple test procedure*, *Scandinavian J. Stat.* 6 (1979), pp. 65–70.
- [14] D.G. Horvitz and D.J. Thompson, *A generalization of sampling without replacement from a finite universe*, *J. Am. Stat. Assoc.* 47 (1952), pp. 663–685.

- [15] M. Huber, M. Lechner and C. Wunsch, *The performance of estimators based on the propensity score*, J. Econom. 175 (2013), pp. 1–21.
- [16] J.D. Kang and J.L. Schafer, *Demystifying double robustness: a comparison of alternative strategies for estimating a population mean from incomplete data*, Stat. Sci. 22 (2007), pp. 523–539.
- [17] E.H. Kennedy, *Semiparametric theory and empirical processes in causal inference*, in *Statistical Causal Inferences and Their Applications in Public Health Research*, H. He, P. Wu, D.-G. Chen, eds., Springer, Cham, 2016, pp. 141–167.
- [18] C. Kooperberg, *polspline: Polynomial Spline Routines*, R package version 1.1.14, 2019. Available at <https://CRAN.R-project.org/package=polspline>.
- [19] C. Kooperberg, S. Bose and C.J. Stone, *Polychotomous regression*, J. Am. Stat. Assoc. 92 (1997), pp. 117–127.
- [20] N. Kreif, S. Gruber, R. Radice, R. Grieve and J.S. Sekhon, *Evaluating treatment effectiveness under model misspecification: a comparison of targeted maximum likelihood estimation with bias-corrected matching*, Stat. Methods. Med. Res. 25 (2016), pp. 2315–2336.
- [21] L. Li and T. Greene, *A weighting analogue to pair matching in propensity score analysis*, Int. J. Biostat. 9 (2013), pp. 215–234.
- [22] F. Li and F. Li, *Propensity score weighting for causal inference with multiple treatments*, Ann. Appl. Stat. 13 (2019), pp. 2389–2415.
- [23] F. Li, K.L. Morgan and A.M. Zaslavsky, *Balancing covariates via propensity score weighting*, J. Am. Stat. Assoc. 113 (2018), pp. 390–400.
- [24] A. Liaw and M. Wiener, *Classification and regression by randomforest*, R News 2 (2002), pp. 18–22. Available at <https://CRAN.R-project.org/doc/Rnews/>.
- [25] M.J. Lopez and R. Gutman, *Estimation of causal effects with multiple treatments: a review and new ideas*, Stat. Sci. 32 (2017), pp. 432–454.
- [26] J.K. Lunceford and M. Davidian, *Stratification and weighting via the propensity score in estimation of causal treatment effects: a comparative study*, Stat. Med. 23 (2004), pp. 2937–2960.
- [27] M.A. Luque-Fernandez, M. Schomaker, B. Rachet and M.E. Schnitzer, *Targeted maximum likelihood estimation for a binary treatment: A tutorial*, Stat. Med. 37 (2018), pp. 2530–2546.
- [28] H. Mao, L. Li and T. Greene, *Propensity score weighting analysis and treatment effect discovery*, Stat. Methods. Med. Res. 28 (2019), pp. 2439–2454.
- [29] E.E. Moodie, T.S. Richardson and D.A. Stephens, *Demystifying optimal dynamic treatment regimes*, Biometrics 63 (2007), pp. 447–455.
- [30] S.A. Murphy, *Optimal dynamic treatment regimes*, J. R. Stat. Soc.: Ser. B (Stat. Methodol.) 65 (2003), pp. 331–355.
- [31] E. Polley, E. LeDell, C. Kennedy and M. van der Laan, *SuperLearner: Super Learner Prediction*, R package version 2.0-24, 2018. Available at <https://CRAN.R-project.org/package=SuperLearner>.
- [32] J.M. Robins, *Marginal structural models*, in *Proceedings of the Section on Bayesian Statistical Science*, Vol. 1997, American Statistical Association, Alexandria, VA, 1998, pp. 1–10.
- [33] J.M. Robins, *Marginal structural models versus structural nested models as tools for causal inference*, in *Statistical Models in Epidemiology, the Environment, and Clinical Trials*, M. Halloran and B. Donald, eds., Springer, New York, 2000, pp. 95–133.
- [34] J.M. Robins, *Optimal structural nested models for optimal sequential decisions*, *Proceedings of the Second Seattle Symposium in Biostatistics*, D. Lin, P.J. Heagerty, eds., Springer, New York, 2004, pp. 189–326.
- [35] A.D. Scotina, F.L. Beaudoin and R. Gutman, *Matching estimators for causal effects of multiple treatments*, Stat. Methods. Med. Res. 29 (2019), pp. 1051–1066.
- [36] A. A. Siddique, M. E Schnitzer, A. Bahamyirou, G. Wang, T. H Holtz, G. B Migliori, G. Sotgiu, N. R Gandhi, M. H Vargas, D. Menzies and A. Benedetti, *Causal inference with multiple concurrent medications: a comparison of methods and an application in multidrug-resistant tuberculosis*, Stat. Methods. Med. Res. 28 (2019), pp. 3534–3549.
- [37] P-O. Tremblay, T. Duchesne, S. G. Cumming and J. A. Jones, *Survival analysis and classification methods for forest fire size*, PLoS One 13 (2018), p. e0189860.
- [38] M.J. Van der Laan, E.C. Polley and A.E. Hubbard, *Super learner*, Stat. Appl. Genet. Mol. Biol. 6 (2007), Article 25.

- [39] M.J. Van der Laan and S. Rose, *Targeted Learning: Causal Inference for Observational and Experimental Data*, Springer, New York, 2011.
- [40] M.J. Van Der Laan and D. Rubin, *Targeted maximum likelihood learning*, Int. J. Biostat. 2 (2006), p. 28.
- [41] Wildfire Management Branch – Alberta Agriculture and Forestry, *Historical Wildfire Database*, 2015. Available at <http://wildfire.alberta.ca/resources/historical-data/historical-wildfire-database.aspx>.
- [42] S. Yang, G.W. Imbens, Z. Cui, D.E. Faries and Z. Kadziola, *Propensity score matching and subclassification in observational studies with multi-level treatments*, Biometrics 72 (2016), pp. 1055–1065.
- [43] K. Yoshida, S. Hernández-Díaz, D.H. Solomon, J.W. Jackson, J.J. Gagne, R.J. Glynn and J.M. Franklin, *Matching weights to simultaneously compare three treatment groups: comparison to three-way matching*, Epidemiology 28 (2017), pp. 387–395.
- [44] G. Zou, *Assessment of risks by predicting counterfactuals*, Stat. Med. 28 (2009), pp. 3761–3781.