



Estimating optimal decision trees for treatment assignment: The case of $K > 2$ treatment alternatives

Aniek Sies¹ · Lisa Doove¹ · Kristof Meers¹ · Elise Dusseldorp² · Iven Van Mechelen¹

Accepted: 26 June 2024
© The Psychonomic Society, Inc. 2024

Abstract

For many problems in clinical practice, multiple treatment alternatives are available. Given data from a randomized controlled trial or an observational study, an important challenge is to estimate an optimal decision rule that specifies for each client the most effective treatment alternative, given his or her pattern of pretreatment characteristics. In the present paper we will look for such a rule within the insightful family of classification trees. Unfortunately, however, there is dearth of readily accessible software tools for optimal decision tree estimation in the case of more than two treatment alternatives. Moreover, this primary tree estimation problem is also cursed with two secondary problems: a structural missingness in typical studies on treatment evaluation (because every individual is assigned to a single treatment alternative only), and a major issue of replicability. In this paper we propose solutions for both the primary and the secondary problems at stake. We evaluate the proposed solution in a simulation study, and illustrate with an application on the search for an optimal tree-based treatment regime in a randomized controlled trial on $K=3$ different types of aftercare for younger women with early-stage breast cancer. We conclude by arguing that the proposed solutions may have relevance for several other classification problems inside and outside the domain of optimal treatment assignment.

Keywords Classification trees · Optimal treatment decisions · Replicability

In clinical practice, often multiple treatment alternatives are available for a problem at hand. In such cases, there is a clear need for decision rules that indicate which treatment alternative should ideally be administered to each client or patient under study. Retrieving such optimal decision rules (also referred to as optimal treatment regimes) is a key methodological challenge in the most topical areas of precision health and precision medicine (Chakraborty & Moodie, 2013; Chakraborty & Murphy, 2014; Huang et al., 2019; Kosorok & Moodie, 2015; Laber et al., 2014; Lou et al., 2018; Schulte et al., 2014; Trivedi, 2016; Whitcomb, 2019).

For the estimation of optimal treatment regimes [limited here to single-decision regimes (Tsiatis et al., 2020), while, for the time being, leaving multiple-decision regimes in a multistage treatment context aside], one has to rely on data

for a sample of clients, each of whom underwent one of the treatment alternatives, within the context of either a randomized clinical trial or an observational study. In the present paper we will reanalyze, as a motivating running example, data from a randomized clinical trial on three types of aftercare administered to 224 younger women with early-stage breast cancer (the Breast Cancer Recovery Project: Scheier et al., 2005, 2007). Prior to the aftercare, the majority of the women underwent a lumpectomy, removal of axillary nodes, and combined radiation and chemotherapy. The three alternative types of aftercare were (1) standard medical care, (2) standard medical care plus a nutrition intervention (how to adopt a low-fat, high-fruit/vegetable diet), and (3) standard medical care plus an education intervention (with information on breast cancer and training of coping skills). Before starting aftercare, 11 pretreatment characteristics were measured. We further focus here on improvement in physical functioning as outcome variable. The key question to be addressed then is which types of women (in terms of pretreatment characteristics) would improve most with regard to physical functioning from which type of aftercare.

✉ Iven Van Mechelen
Iven.VanMechelen@kuleuven.be

¹ University of Leuven, Tiensestraat 102 Box 3713,
3000 Leuven, Belgium

² Leiden University, Leiden, The Netherlands

When looking for an optimal treatment regime, one typically must do so within a prespecified family or class of treatment regimes. In the present paper we will focus for this purpose on the family of classification trees, because of their obvious insightfulness. (For examples of classification trees, see Figs. 3 and 4.) At this point, however, an obstacle arises: The problem of estimating an optimal decision tree-based treatment regime has been satisfactorily analyzed at a theoretical level (Laber & Zhao, 2015; Tsiatis et al., 2020; Zhang et al., 2012); a ready-made and easily accessible software solution has further been made available for the case of two treatment alternatives (Holloway et al., 2023; Zhang et al., 2012); however, a similar readily accessible solution for the case of more than two treatment alternatives (as in the running example outlined above) is not yet available (with, e.g., Laber & Zhao, 2015, sharing only a very partial and limited piece of pseudocode in supplementary materials). In the present paper, we will solve for this.

The remainder of this paper is structured as follows: In Sect. 2 we will first formalize and analyze the primary problem that is the focus of the present paper (i.e., the lack of a readily accessible methodology for estimating optimal tree-based treatment regimes in the case of strictly more than two treatment alternatives); subsequently, we will propose a solution for it. Second, we will also briefly discuss two secondary problems that arise during the estimation, along with a proposed solution for them. Section 3 will present an evaluation of the proposed methodology in a simulation study, and Sect. 4 an application of it to the data from the Breast Cancer Recovery Project. We will end with a few concluding remarks.

Method

Primary problem

Notation and formalization

The methodology we propose in the present paper is applicable to data from both randomized clinical trials (RCT) and observational studies. For simplicity's sake, however, we will chiefly zoom in on the RCT case.

We therefore assume that data are available from an RCT that involves I individuals, clients or patients, $(1, \dots, i, \dots, I)$. In the running example of the Breast Cancer Recovery Project, $I = 224$.

In the RCT, each individual is randomly assigned to one out of K treatment alternatives $(1, \dots, k, \dots, K)$. In the running example, $K = 3$. We further denote the variable that indicates the treatment alternative to which each individual is assigned in the RCT by A .

From each individual, a set of P pretreatment or baseline characteristics \mathbf{X} is collected prior to treatment assignment, with $\mathbf{X} = (X_1, \dots, X_p, \dots, X_P)$. In the running example, $P = 11$. In addition, an outcome variable Y is measured; without loss of generality, we further assume that higher values on Y are better.

A treatment regime g then can be defined as a mapping:

$$g : \text{Range}(\mathbf{X}) \rightarrow \{1, \dots, k, \dots, K\} \\ \mathbf{x} \mapsto g(\mathbf{x})$$

Note that this definition also includes constant functions as treatment regimes—that is to say, so-called trivial or one-size-fits-all treatment regimes, which imply that all individuals are assigned to one and the same treatment alternative.

To formalize treatment regime optimality, we firstly have to prespecify a family of treatment regimes \mathcal{G} (with most families also including the one-size-fits-all regimes as special cases). In the present paper, \mathcal{G} is the family of classification trees. Secondly, we need the concept of potential outcomes (Rubin, 1974), with, in the case of K treatment alternatives, K potential outcome variables, $\mathbf{Y}^* = (Y^{*1}, \dots, Y^{*k}, \dots, Y^{*K})$, where Y^{*k} denotes the outcome that would have been observed if the individual under study were assigned to treatment alternative k . (One may note that, as each individual is assigned to a single treatment alternative, effectively only one of the potential outcome variables is observed—which comes down to a problem of structural missingness; we will return to this issue in the discussion of the secondary problems.) Thirdly, we need an optimality criterion to define within the search space \mathcal{G} an optimal treatment regime g^{opt} . Several criteria are possible in this regard, with the criterion most often used in practice being that of maximizing the expected potential outcome (Tsiatis et al., 2020). If we denote by $Y^{*g(\mathbf{X})}$ the random variable that takes for individual i the value $Y^{*g(\mathbf{X}(i))}$, this criterion can be written as

$$g^{\text{opt}} = \underset{g \in \mathcal{G}}{\text{argmax}} E[Y^{*g(\mathbf{X})}] = \underset{g \in \mathcal{G}}{\text{argmax}} E_{\mathbf{X}}[E[Y^{*g(\mathbf{X})} | \mathbf{X}]]. \quad (1)$$

Analysis of primary problem

We will now analyze the primary problem as formalized by Eq. (1), with \mathcal{G} being the family of classification trees. We will do so while generalizing derivations by Zhang et al. (2012) and Tsiatis et al. (2020). Ultimately, our analysis will result in a (counterintuitive) transformation of our primary optimal treatment regime estimation problem into a supervised classification problem, with the truly optimal treatment alternative for each client acting as the “supervisor” (i.e., as that client’s “true class”). This is counterintuitive, indeed, as the truly optimal treatment alternative for each client is

typically unknown. The key to this counterintuitive mystery is that, in an initial, preparatory stage, the true class for each client (as well as the so-called client-specific misclassification costs that will be further explained below) will be estimated on the basis of the RCT data (including the observed treatment assignment in that RCT). The latter will be further discussed below as the first of the secondary problems.

For our analysis, we first denote

$$E[Y|A = a, \mathbf{X}] = \mu(a, \mathbf{X}).$$

We further assume (Rubin, 2005) that $Y(i) = Y^{*A(i)}(i)$. If we also assume independence of \mathbf{Y}^* and A conditional on \mathbf{X} (which is trivial in case of an RCT and requires a so-called assumption of no unmeasured confounders in observational studies), it follows that

$$\mu(a, \mathbf{X}) = E[Y^{*a}|\mathbf{X}].$$

In that case, (1) can be rewritten as

$$g^{opt} = \operatorname{argmax}_{g \in G} E_{\mathbf{X}}[\mu(g(\mathbf{X}), \mathbf{X})].$$

This further implies that

$$\begin{aligned} g^{opt} &= \operatorname{argmin}_{g \in G} E_{\mathbf{X}} \left[\max_k \mu(k, \mathbf{X}) - \mu(g(\mathbf{X}), \mathbf{X}) \right] \\ &= \operatorname{argmin}_{g \in G} E_{\mathbf{X}} \left[1_{g(\mathbf{X}) \neq \operatorname{argmax}_k \mu(k, \mathbf{X})} \left(\max_k \mu(k, \mathbf{X}) - \mu(g(\mathbf{X}), \mathbf{X}) \right) \right]. \end{aligned}$$

Hence, assuming that for each client i , (estimates of) all $\mu(k, \mathbf{x}_i)$ values are available (and, hence, also the values of $\max_k \mu(k, \mathbf{x}_i)$ and $\operatorname{argmax}_k \mu(k, \mathbf{x}_i)$), on the sample level a classification tree g is to be looked for that minimizes

$$\frac{1}{I} \sum_i 1_{g(\mathbf{x}_i) \neq \operatorname{argmax}_k \mu(k, \mathbf{x}_i)} \left[\max_k \mu(k, \mathbf{x}_i) - \mu(g(\mathbf{x}_i), \mathbf{x}_i) \right]. \quad (2)$$

Such a tree is typically built by starting from a root node that includes all units (individuals), and by subsequently recursively partitioning each of the end nodes of the current tree on the basis of a covariate and a split point that are chosen such as to minimize (2).

To arrive at a deeper understanding of Eq. (2), the interpretations listed in Table 1, which involve a transition from

“individual” and “treatment alternative” to the more generic terms of “unit” and “class” that are customary in the classification domain, may be helpful. These interpretations imply that Eq. (2) comes down to a search for a classification tree that minimizes an average or total misclassification cost.

To further put this in a broader context, we may consider four different possible types of misclassification costs (summarized in Table 2) that are distinguished in the classification domain (see, e.g., Höppner et al., 2022, and references therein), the fourth of which applies to our primary problem.

The first line of this table corresponds to what Feng et al. (2021) call the “classical classification paradigm,” which aims to minimize the overall misclassification rate. In contrast, the lower three lines correspond to three types of “cost-sensitive learning.” In particular, the class-dependent type implies that the misclassification cost may vary across pairs of true and assigned classes. Note that the index i in $T(i)$ implies unit dependency only via the true class of unit i , with the misclassification costs being fully captured by a $K \times K$ matrix. Note further that the latter matrix may be asymmetric, as in the example of classification of email messages as non-spam versus spam, where erroneously misclassifying truly non-spam mail as spam may be more costly than the other way around. As an example of unit- (instance-, example-, or case-)dependent costs, one may think of credit scoring, where the cost of misclassifying a customer as creditworthy (vs. not creditworthy) may depend on the amount that the customer in question wants to borrow.

In our case, Eq. (2) implies that the misclassification cost depends both on the client (unit) and on the treatment alternative (class) to which the client is (mis)classified by the treatment regime. Hence, we are dealing with a unit- and class-dependent misclassification cost, indeed.

Table 2 Types and forms of misclassification costs

Misclassification cost	Type	Form
$m_{ic} = mcc$	Constant	1×1 scalar
$m_{ic} = mcc_{T(i)c}$	Class-dependent	$K \times K$ matrix
$m_{ic} = mcc_i$	Unit-dependent	$I \times 1$ vector
$m_{ic} = mcc_{ic}$	Unit- and class-dependent	$I \times K$ matrix

Table 1 Interpretation of terms of Eq. (2)

Term in Eq. (2)	Interpretation
$g(\mathbf{x}_i) = c$	Class (treatment alternative) to which unit (individual) i is assigned according to classification tree g
$\operatorname{argmax}_k \mu(k, \mathbf{x}_i) = T(i)$	True class (optimal treatment alternative) for unit (individual) i
$1_{g(\mathbf{x}_i) \neq \operatorname{argmax}_k \mu(k, \mathbf{x}_i)}$	Whether or not unit (individual) i is misclassified by classification tree g
$[\max_k \mu(k, \mathbf{x}_i) - \mu(g(\mathbf{x}_i), \mathbf{x}_i)] = m_{ic}$	Misclassification cost implied by erroneously assigning unit (individual) i to class (treatment alternative) c instead of to its true class (optimal treatment alternative) $T(i)$

Solution to the primary problem

In the search for a readily accessible software tool for estimating classification trees with minimal misclassification cost, an obvious choice is the mainstream package for tree estimation *rpart* (Therneau et al., 2022). The *rpart* package can deal with misclassification costs of different types: A constant cost is the default option, the class-dependent case can be dealt with by making use of the option of a loss matrix, and the unit-dependent case can be handled by using the option of case weights. For the unit- and class-dependent case, however, the situation is somewhat different. If $K = 2$ and if correct classifications go with a zero cost, one can make use of the fact that the $I \times 2$ misclassification cost matrix can be reduced to a vector of length I ; the latter can be included in *rpart* via the weight vector. If $K > 2$, however, so far no solution has been made available.

We propose for this purpose a novel code snippet *multivalued*, a fully documented version of which (including all code used in the analysis of the running example, as reported below) is available through GitHub (<https://github.com/KULeuven-PPW-OKPIV/multivalued>), and a significant part of which is also available in the online supplementary material. This code snippet is to be used in conjunction with *rpart*, and is based on the following building blocks:

1. The option in *rpart* to make use of a user-defined splitting function (Therneau, 2022)
2. The possibility to enter in the optional cost matrix in *rpart* a rectangular $I \times K$ matrix
3. The possibility to enter in the weight vector of *rpart* the unit numbers, which subsequently can be used to select for each unit the relevant row from the cost matrix

Note that the user-defined splitting function implemented in the code snippet implies that in the tree building, the total misclassification cost as formalized by Eq. (2) is directly minimized (rather than the Gini index that is used as default criterion in *rpart*). Note further that, whereas utilizing a user-defined split function in *rpart* can slow down the package's execution (Therneau, 2022), in our case the slowing down appeared not to be prohibitively large (with, e.g., each of the analyses in our illustrative application taking less than one second).

Secondary problems

Two secondary problems show up in the estimation of tree-based optimal treatment regimes and require a satisfactory solution. The first of these pertains to the estimation of the true class and the misclassification costs for each client. This is hampered by the structural missingness in typical observational studies and RCTs on treatment evaluation, which

is because every individual is assigned to a single treatment alternative only. This implies that for each individual, a single potential outcome is known. In the literature, various solutions have been proposed to deal with this issue (Laber & Zhao, 2015; Tsiatis et al., 2020; Zhang et al., 2012). A first of these comes down to (the estimation of) a regression-based outcome model, in which the outcome is modeled as a function of baseline characteristics, treatment assignment, and the interaction between them (Tsiatis et al., 2020). This approach (also called Q-learning, with "Q" referring to "quality"), though, strongly depends on whether the outcome model in question has been correctly specified. As a way out, one may consider relying on flexible modeling solutions, such as random forests (Laber & Zhao, 2015). Still another possibility is to revert to more robust estimators of potential outcomes and their contrasts (Tsiatis et al., 2020; Zhang et al., 2012); such estimators include inverse probability weighted estimators (IPWE, with the probabilities in question pertaining to propensity models to deal with observational studies), and augmented inverse probability weighted estimators (AIPWE, with the augmentation pertaining to a term stemming from an outcome model). A particularly promising characteristic of AIPWEs is their so-called double robustness, meaning that they can be shown to be consistent if the propensity model *or* the outcome model has been correctly specified (with the propensity model being trivially correct in the case of an RCT).

A second secondary problem to be dealt with is the issue of replicability. This is an issue of concern because of the replicability crisis in many disciplines of science, including not least in psychology (see, e.g., Klein et al., 2018). Moreover, the issue is of particular relevance in the estimation of treatment regimes, and more generally in exploratory subgroup analyses in clinical applications, as in the past such analyses have not infrequently led to conclusions that could not be replicated (e.g., Rothwell, 2005). As a consequence, some authors even disposed of exploratory subgroup analyses as "data dredging" (Feinstein, 1998; Rothwell, 2005). This replicability problem relates in part to the structural missingness referred to above and the estimation uncertainty implied by it. Furthermore, it also immediately relates to the fact that, compared to main effects, considerably larger sample sizes are needed for a reliable estimation of treatment by subgroup interactions, and in particular for the detection of qualitative or disordinal such interactions that constitute the basis of nontrivial treatment regimes (Brookes et al., 2004). On top of all this comes the long-known problem of instabilities of trees with regard to variables and split points (Breiman et al., 1984). (One might argue that the mainstream package for tree estimation *rpart* includes a cross-validation procedure, which could be considered as a kind of protection to safeguard replicability; yet, in the package, cross-validation is used only in a pruning procedure after the tree building, and as such

only yields a limited way out for the instabilities issue.) As a solution for all these problems, we propose making an appeal to multiverse analysis (Steege et al., 2016), that is to say, to make use of a broad range of analysis options, and to subsequently check which results consistently emerge across this range.

Simulation study

We set up a Monte Carlo simulation study to evaluate the performance of three variants of the methodology to estimate optimal tree-based treatment regimes in RCTs for which we proposed an accessible software solution. The three variants are based on three different ways to estimate the true class and the misclassification costs for each client: (1) a flexible outcome model based on random forests (more information on which will be provided in the Application section), (2) a classical AIPWE approach based on a linear regression outcome model (that includes all baseline covariates and their interaction with treatment) and theoretical propensity probabilities, and (3) the same AIPWE approach in which the theoretical propensity probabilities are replaced by empirical proportions (which, somewhat paradoxically, has been shown to yield a more efficient estimation: see Tsiatis, 2006, p. 206; Tsiatis et al., 2020, p. 44–46). We wanted to evaluate the performance of these methods across a range of settings that varied in terms of data characteristics that could affect that performance. In particular, we generated data $(Y_i, A_i, X_{1i}, \dots, X_{5i})$, $i = 1, \dots, I$, with (X_1, \dots, X_5) iid standard normal, A multinomial with $P(A = a_k) = \frac{1}{K}$, and Y generated according to

$$Y_i = 1.0 + 0.25X_{1i} + 0.25X_{2i} - 0.25X_{5i} - \theta 1_{a_i \neq g^{opt}(\mathbf{x}_i)} + E_i,$$

with θ a prespecified constant, $g^{opt}(\mathbf{x}_i)$ the truly optimal treatment alternative for individual i as defined below, and E_i standard normal. In the data generation, we further systematically varied the following four data characteristics in a full factorial design (with 500 data sets within each cell):

- (1) The number of arms in the RCT, $K = 3, 4$
- (2) The sample size per arm, $n_K = 50, 100, 200$
- (3) The effect size of the difference in outcome between optimal and non-optimal treatment alternatives in each relevant subgroup, $\theta = 0.50, 1$
- (4) The true optimal treatment regime (OTR) underlying the data, based on three scenarios—a tree-based, a non-tree-based, and a one-size-fits-all or trivial one—as depicted in Fig. 1

We subjected each simulated data set to each of the three OTR estimation methods under study in conjunction with *rpart*. For the pruning of the estimated optimal tree-based treatment regimes, we used 20-fold cross-validation.

We focused on four performance aspects:

- (1) The expected outcome of each estimated regime, in terms of its normalized performance gain (NPG), which compares the benefit gained by administering that regime over administering the marginally best treatment alternative a^{opt} to the benefit that could have been gained theoretically,

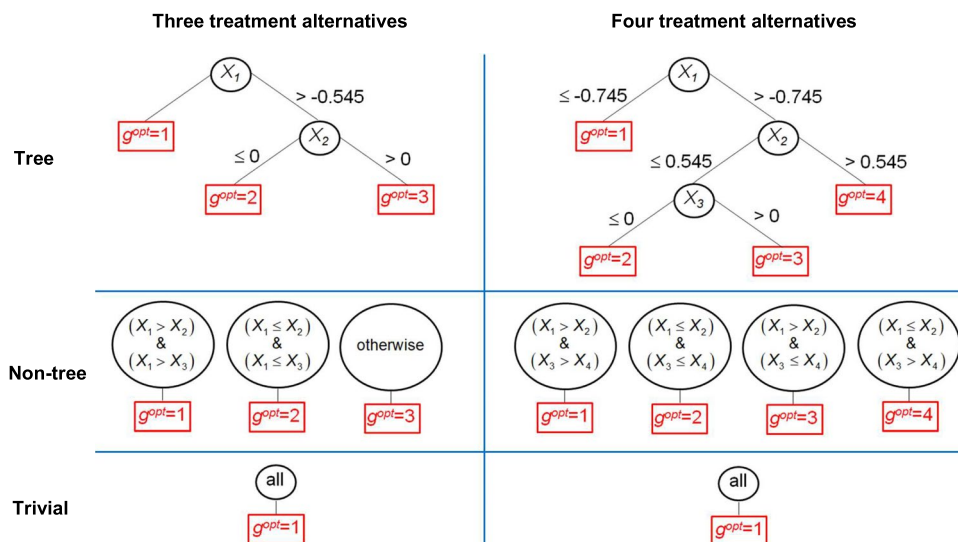


Fig. 1 True optimal treatment regimes used in simulation study with tree-based (row 1), non-tree-based (row 2), and one-size-fits-all or trivial (row 3) scenario and with three (column 1) and four (column 2) treatment alternatives

$$NPG = \frac{E(Y^{*\hat{g}^{opt}}) - E(Y^{*\alpha^{opt}})}{E(Y^{*g^{opt}}) - E(Y^{*\alpha^{opt}})},$$

where $E(Y^{*\hat{g}^{opt}})$ and $E(Y^{*\alpha^{opt}})$ were calculated on the basis of a simulated “super-sample” of 10^6 observations generated on the basis of the true model underlying the simulated data.

- (2) Classification accuracy, that is, the proportion of patients assigned to their truly optimal treatment alternative.
- (3) The “Type I error rate” (with a slight abuse of terminology), that is, the proportion of data sets in a cell generated under a one-size-fits-all scenario for which a method erroneously yielded a nontrivial estimated OTR.
- (4) The “Type II error rate” (with again a slight abuse of terminology), that is, the proportion of data sets in a cell generated under a nontrivial scenario for which a method erroneously yielded a one-size-fits-all estimated OTR.

To identify the most important effects, we subjected the four outcome measures to a repeated-measures ANOVA, with the within-factor pertaining to the type of OTR estimation method and the between-factors to the data characteristics. We will further only discuss effects with an effect size $\eta^2 \geq 0.05$ (see Table 3).

Sizable main effects were primarily found, the contents of which can be derived from Table 4. As could have been expected, a better performance (in terms of expected outcome and classification accuracy) was found for tree-based as compared to non-tree-based scenarios, for a larger effect size of the differences in outcome between optimal and non-optimal treatment alternatives, and for larger sample sizes. Furthermore, inferential error rates appeared to be better as well in the case of a larger effect size of the differences in outcome between optimal and non-optimal treatment alternatives, and of larger sample sizes; in addition, the random forest-based method appeared to yield slightly better inferential results, with for the “Type I error rate” this being

especially true in the case of smaller sample sizes per treatment arm (Fig. 2).

Illustrative application

Analysis

The data of the Breast Cancer Recovery Project are publicly available as part of the R package *quint* (Dusseldorp et al., 2016, 2022). As already indicated above, all code for the analyses described in the present section is available through GitHub (<https://github.com/KULeuven-PPW-OKPIV/multi-valued>) and in the online supplementary material.

We analyze the Breast Cancer Recovery data, with all 11 pretreatment characteristics measured at baseline as covariates, including age, physical functioning, and number of comorbidities (such as diabetes, migraine, arthritis, angina, ...), and with physical functioning measured at 9-month follow-up minus physical functioning measured at baseline (i.e., improvement in physical functioning) as outcome variable; note that possible objections against the use of a change score as an outcome variable (e.g., Senn, 2006) do not apply here, as we also included physical functioning at baseline as covariate. We estimate the misclassification costs via, on the one hand, a flexible outcome modeling technique, that is to say, random forests, and, on the other hand, AIPWE.

For the random forests, we make use of the *ranger* package (Wright & Ziegler, 2017), with default values for the tuning parameters (including 500 trees per forest, and with the square root of the total number of variables as the size of the random subset of variables that are considered for the splits in each tree). Moreover, we impose the constraint that the treatment assignment variable is always a member of the subset of variables considered for splits, to avoid a bias of the random forest towards trivial one-size-fits-all regimes (which would be implied by a random non-inclusion of treatment assignment). Importantly, the constraint mentioned above is that treatment assignment is always *considered* for the splits; therefore, it always leaves room for trees *without* treatment

Table 3 Effect size (η^2) of effects in simulation study with $\eta^2 \geq .05$

Effect	Outcome variable			
	Normalized performance gain	Classification accuracy	“Type I error rate”	“Type II error rate”
Sample size per arm	.16	.05	.07	.33
Effect size	.20	.08	.66	.35
Scenario	.23	.59		.07
Method			.11	.05
Method * Sample size per arm			.05	

Table 4 Marginal means on four outcome variables for categories of five design variables in simulation study. Marginal means for main effects with effect size $\eta^2 \geq .05$ are depicted in bold

Effect	Outcome measure			
	Normalized performance gain	Classification accuracy	“Type I error rate”	“Type II error rate”
Number of arms				
3	.43	.75	.17	.08
4	.32	.67	.15	.10
Sample size per arm				
50	.23	.64	.20	.16
100	.38	.71	.15	.08
200	.53	.79	.13	.03
Effect size				
0.5	.24	.64	.24	.14
1	.52	.79	.07	.04
Scenario				
Tree	.53	.68	NA	.07
Non-tree	.23	.48	NA	.11
Trivial	NA	.97	.16	NA
Method				
AIPWE empirical proportions	.37	.71	.17	.11
AIPWE probabilities	.39	.71	.19	.09
Random Forest	.37	.72	.11	.06

assignment as splitting variable and, hence, without treatment effect heterogeneity. Ultimately, it will be up to the data to decide whether or not treatment effect heterogeneity will be in place in the random forest. This further implies that the constrained random forest approach can perfectly lead to one-size-fits-all estimated optimal treatment regimes (as will also appear below).

For the AIPWE, we use for the propensity part both the theoretical probabilities and the empirical proportions of individuals assigned to the three treatment arms (standard medical care: 0.34, nutrition: 0.35, education: 0.31); for the outcome model part, we use a regression model that

includes as predictors all baseline characteristics as well as their interactions with treatment assignment.

For the multiverse analysis, in addition to using both random forests and AIPWE for the misclassification cost estimation, and in addition to the use of both theoretical probabilities and empirical proportions for the inverse probability weighing, we use for the random forests three random starts; furthermore, we repeat the 20-fold cross-validation five times for the pruning in *rpart*. All this results in 3 (random forest start) \times 5 (pruning cross-validation) = 15 random forest-based solutions, and 2 (probabilities vs. empirical proportions) \times 5 (pruning cross-validation) = 10 AIPWE-based solutions.




Fig. 2 Method by sample size per arm interaction for outcome measure “Type I error rate” in simulation study

Results

Out of the 15 random forest-based solutions, nine solutions are a tree with a root only, that is to say, a trivial, one-size-fits-all treatment regime in which all individuals are assigned to education; furthermore, five solutions are a nontrivial regime that involves the baseline characteristics age, number of comorbidities, and physical functioning at baseline. This regime is graphically represented in Fig. 3. It implies that relatively younger women with a higher number of comorbidities (such as diabetes, migraine, etc.) and a relatively higher level of physical functioning at baseline should

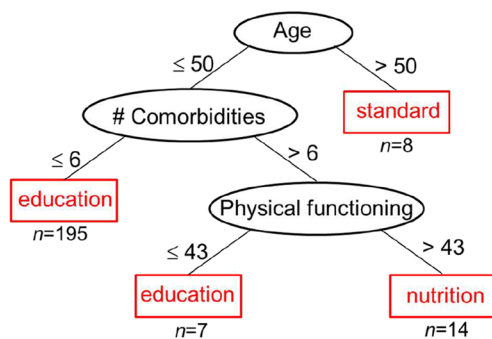


Fig. 3 Nontrivial random forest-based optimal decision tree for Breast Cancer Recovery data

ideally receive the nutrition intervention, whereas relatively older women should ideally receive standard medical care. Finally, one nontrivial solution involves age only.

Out of the 10 AIPWE-based solutions, 6 solutions pertain to a trivial, one-size-fits-all treatment regime in which all individuals are assigned to education; 2 solutions pertain to the nontrivial regime that is graphically represented in Fig. 4 and that involves the number of comorbidities once as the only covariate (with women with a relatively higher number of comorbidities ideally being assigned to the nutrition intervention); 1 nontrivial solution further involves the number of comorbidities twice and 1 nontrivial solution involves the number of comorbidities twice in addition to age.

Discussion

The whole of our analyses yielded only fairly weak evidence for a nontrivial optimal treatment regime, with 60% of both the random forest- and the AIPWE-based solutions being of the one-size-fits all type, and with the education intervention overall seeming the most beneficial. That being said, 9 out of the 25 obtained solutions included some indication that clients who suffer from a higher number of comorbidities (diabetes, migraine, arthritis, angina, etc.) might benefit more from a nutrition intervention.




Fig. 4 Nontrivial AIPWE-based optimal decision tree for Breast Cancer Recovery data

With about 75 clients per arm, this RCT can be situated between the first and the second sample size levels of the simulation study in the previous section. This goes with somewhat higher “Type I” and “Type II” error rates, which further adds to a somewhat higher inferential uncertainty.

Concluding remarks

In this paper we proposed a readily accessible methodology for estimating optimal classification trees that minimize a loss function involving a unit- and class-based misclassification cost in the case of a classification problem with $K > 2$ classes. In conjunction with a multiverse approach to both misclassification cost and tree estimation, the proposed methodology provided an insightful but shaded solution to a problem of optimal treatment assignment with $K > 2$ treatment alternatives.

We focused here on the use of our proposed methodology within the context of RCTs. The methodology, however, and in particular the AIPWE variant with propensity probabilities, is in principle also applicable to data from observational studies. To be sure, this type of application requires the (untestable) assumption that there are no unmeasured covariates that are associated with both treatment assignment and the potential outcomes (“no unmeasured confounders”: Tsiatis et al., 2020, p. 27ff.)—in addition to other assumptions such as the requirement that the probability for assignment to each treatment alternative be strictly positive in each area of the covariate space.

The proposed methodology was introduced as a tool to estimate optimal tree-based single-decision regimes that involve $K > 2$ treatment alternatives. However, in line with the suggestion made by Holloway et al. (2023) for their $K = 2$ method, by calling our proposed methodology repeatedly, its use can also be extended to multiple-decision regimes in a multistage treatment context (with the additional option to include in that case at some decision point also all evolving client information available at that point as covariates).

The proposed methodology was further introduced as a tool to address treatment-related decisions in a clinical context. That being said, though, the methodology could also be used to address formally similar problems in different areas. As a first example, one may think of the choice between $K > 2$ types of churn management in customer retention (Lemmens & Gupta, 2020), which is a context where misclassification cost can be taken literally and expressed in monetary terms. As a second example, one may think of an identification of optimal data-analytic regimes in data-analytic benchmarking that involves a comparison of $K > 2$ data-analytic methods (Doove et al., 2017); otherwise, the latter type of application is technically simpler, as in a benchmarking context the outcome of all methods under study when applied to a data set at hand is typically known.

Finally, multiverse approaches similar to the ones that we proposed and applied to cope both with replicability challenges in optimal treatment regime estimation and with stability challenges in tree-based analyses, may be applied more broadly as well. As examples one may think of optimal tree-based treatment regime estimation in the case of two treatment alternatives, of optimal treatment regime searches within families of regimes other than the tree-based ones, and of tree analyses based on structures other than simple classification trees, such as regression trees and model-based recursive partitioning (Zeileis et al., 2008).

Supplementary information The online version contains supplementary material available at <https://doi.org/10.3758/s13428-024-02470-9>.

Acknowledgement The authors thank Geert Verbeke for helpful comments on an earlier version of this paper.

Funding The research reported in this paper was supported in part by the Research Foundation—Flanders (G080219N) and by the Research Fund of KU Leuven (C14/19/054).

Data availability The data of the Breast Cancer Recovery Project are publicly available as part of the R package quint (Dusseldorp et al., 2016, 2022).

Code availability All code for the simulation study and the analyses described in the present paper is available through GitHub (<https://github.com/KULeuven-PPW-OKPIV/multivalued>), with the code for reproducing the illustrative application also available in the online supplementary material.

Declarations

Conflicts of interest The authors report no conflicts of interest.

Ethics approval Not applicable (secondary data). The original paper (Scheier et al., 2005) reported: “This study was performed in accordance with an assurance filed with and approved by the Department of Health and Human Services and had the approval of relevant institutional review boards.”

Consent to participate Not applicable (secondary data). The original paper (Scheier et al., 2005) reported: “All participants signed an informed consent.”

Consent for publication Not applicable.

References

- Breiman, L., Friedman, J. H., Olshen, R. A., & Stone, C. J. (1984). Classification and regression trees. *Wadsworth*. <https://doi.org/10.1201/9781315139470>
- Brookes, S. T., Whitely, E., Egger, M., Smith, G. D., Mulheran, P. A., & Peters, T. J. (2004). Subgroup analyses in randomized trials: Risks of subgroup-specific analyses; power and sample size for the interaction test. *Journal of Clinical Epidemiology*, *57*, 229–236. <https://doi.org/10.1016/j.jclinepi.2003.08.009>
- Chakraborty, B., & Moodie, E. (2013). Statistical methods for dynamic treatment regimes: Reinforcement learning, causal inference, and personalized medicine. *Springer*. <https://doi.org/10.1007/978-1-4614-7428-9>
- Chakraborty, B., & Murphy, S. A. (2014). Dynamic treatment regimes. *Annual Review of Statistics and Its Application*, *1*(1), 447–464. <https://doi.org/10.1146/annurev-statistics-022513-115553>
- Doove, L. L., Wilderjans, T. F., Calcagni, A., & Van Mechelen, I. (2017). Deriving optimal data-analytic regimes from benchmarking studies. *Computational Statistics and Data Analysis*, *107*, 81–91. <https://doi.org/10.1016/j.csda.2016.10.016>
- Dusseldorp, E., Doove, L., van de Put, J., Ninaber, C., Van Mechelen, I., & Claramunt Gonzalez, J. (2022). *Quint: Qualitative interaction trees*. <https://cran.r-project.org/package=quint>
- Dusseldorp, E., Doove, L., & Van Mechelen, I. (2016). Quint: An R package for the identification of subgroups of clients who differ in which treatment alternative is best for them. *Behavior Research Methods*, *48*, 650–663. <https://doi.org/10.3758/s13428-015-0594-z>
- Feinstein, A. R. (1998). The problem of cogent subgroups: A clinico-statistical tragedy. *Journal of Clinical Epidemiology*, *51*, 297–299. [https://doi.org/10.1016/S0895-4356\(98\)00004-3](https://doi.org/10.1016/S0895-4356(98)00004-3)
- Feng, Y., Zhou, M., & Tong, X. (2021). Imbalanced classification: A paradigm-based review. *Statistical Analysis and Data Mining: The ASA Data Science Journal*, *14*(5), 383–406. <https://doi.org/10.1002/sam.11538>
- Holloway, S. T., Laber, E. B., Linn, K. A., Zhang, B., Davidian, M., & Tsiatis, A. A. (2023). *DynTxRegime: Methods for estimating optimal dynamic treatment regimes*. <https://cran.r-project.org/package=DynTxRegime>
- Höppner, S., Baesens, B., Verbeke, W., & Verdonck, T. (2022). Instance-dependent cost-sensitive learning for detecting transfer fraud. *European Journal of Operational Research*, *297*, 291–300. <https://doi.org/10.1016/j.ejor.2021.05.028>
- Huang, X., Goldberg, Y., & Xu, J. (2019). Multicategory individualized treatment regime using outcome weighted learning. *Biometrics*, *75*(4), 1216–1227. <https://doi.org/10.1111/biom.13084>
- Klein, R. A., Vianello, M., Hasselman, F., Adams, B. G., Adams, R. B. Jr., Alper, S., Aveyard, M., Axt, J. R., Babalola, M. T., Bahník, S., Batra, R., Berkics, M., Bernstein, M. J., Berry, D. R., Bialobrzeska, O., Binan, E. D., Bocian, K., Brandt, M. J., Busching, R., . . . Nosek, B. A. (2018). Many labs 2: Investigating variation in replicability across samples and settings. *Advances in Methods and Practices in Psychological Science*, *1*(4), 443–490. <https://doi.org/10.1177/2515245918810225>
- Kosorok, M., & Moodie, E. (Eds.). (2015). *Adaptive treatment strategies in practice: Planning trials and analyzing data for personalized medicine*. Society for Industrial and Applied Mathematics. <https://doi.org/10.1137/1.9781611974188>
- Laber, E., Lizotte, D. J., Qian, M., Pelham, W. E., & Murphy, S. A. (2014). Dynamic treatment regimes: Technical challenges and applications. *Electronic Journal of Statistics*, *8*(1), 1225–1272. <https://doi.org/10.1214/14-ejs920>
- Laber, E. B., & Zhao, Y. Q. (2015). Tree-based methods for individualized treatment regimes. *Biometrika*, *102*, 501–514. <https://doi.org/10.1093/biomet/asv028>
- Lemmens, A., & Gupta, S. (2020). Managing churn to maximize profits. *Marketing Science*, *39*, 956–973. <https://doi.org/10.1287/mksc.2020.1229>
- Lou, Z., Shao, J., & Yu, M. (2018). Optimal treatment assignment to maximize expected outcome with multiple treatments. *Biometrics*, *74*(2), 506–516. <https://doi.org/10.1111/biom.12811>
- Rothwell, P. M. (2005). Subgroup analysis in randomised controlled trials: Importance, indications, and interpretation. *The Lancet*, *365*, 176–186. [https://doi.org/10.1016/S0140-6736\(05\)17709-5](https://doi.org/10.1016/S0140-6736(05)17709-5)
- Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, *66*, 688–701. <https://doi.org/10.1037/h0037350>

- Rubin, D. B. (2005). Causal inference using potential outcomes: Design, modeling, decisions. *Journal of the American Statistical Association*, *100*(469), 322–331. <https://doi.org/10.1198/016214504000001880>
- Scheier, M. F., Helgeson, V. S., Schulz, R., Colvin, S., Berga, S., Bridges, M. W., Knapp, J., Gerszten, K., & Pappert, W. S. (2005). Interventions to enhance physical and psychological functioning among younger women who are ending nonhormonal adjuvant treatment for early stage breast cancer. *Journal of Clinical Oncology*, *23*, 4298–4311. <https://doi.org/10.1200/JCO.2005.05.362>
- Scheier, M. F., Helgeson, V. S., Schulz, R., Colvin, S., Berga, S. L., Knapp, J., & Gerszten, K. (2007). Moderators of interventions designed to enhance physical and psychological functioning among younger women with early-stage breast cancer. *Journal of Clinical Oncology*, *25*, 5710–5714. <https://doi.org/10.1200/JCO.2007.11.7093>
- Schulte, P. J., Tsiatis, A., Laber, E., & Davidian, M. (2014). Q-and A-learning methods for estimating optimal dynamic treatment regimes. *Statistical Science*, *29*(4), 640–661. <https://doi.org/10.1214/13-STS450>
- Senn, S. (2006). Change from baseline and analysis of covariance revisited. *Statistics in Medicine*, *25*(24), 4334–4344. <https://doi.org/10.1002/sim.2682>
- Steenen, S., Tuerlinckx, F., Gelman, A., & Vanpaemel, W. (2016). Increasing transparency through a multiverse analysis. *Perspectives on Psychological Science*, *11*, 702–712. <https://doi.org/10.1177/1745691616658637>
- Therneau, T. (2022). *User written splitting functions for rpart*. <https://cran.r-project.org/web/packages/rpart/vignettes/usercode.pdf>
- Therneau, T., Atkinson, B., & Ripley, B. (2022). *Package rpart: Recursive partitioning and regression trees*. <https://cran.r-project.org/package=rpart>
- Trivedi, M. H. (2016). Right patient, right treatment, right time: Biosignatures and precision medicine in depression. *World Psychiatry*, *15*(3), 237–238. <https://doi.org/10.1002/wps.20371>
- Tsiatis, A. A. (2006). Semiparametric theory and missing data. *Springer*. <https://doi.org/10.1007/0-387-37345-4>
- Tsiatis, A. A., Davidian, M., Holloway, S. T., & Laber, E. T. (2020). Dynamic treatment regimes: Statistical methods for precision medicine. *CRC Press*. <https://doi.org/10.1201/9780429192692>
- Whitcomb, D. C. (2019). Primer on precision medicine for complex chronic disorders. *Clinical and Translational Gastroenterology*, *10*(7), e00067–e00067. <https://doi.org/10.14309/ctg.0000000000000067>
- Wright, M. N., & Ziegler, A. (2017). Ranger: A fast implementation of random forests for high dimensional data in C++ and R. *Journal of Statistical Software*, *77*(1), 1–17. <https://doi.org/10.18637/jss.v077.i01>
- Zeileis, A., Hothorn, T., & Hornik, K. (2008). Model-based recursive partitioning. *Journal of Computational and Graphical Statistics*, *17*(2), 492–514. <https://doi.org/10.1198/106186008X319331>
- Zhang, B., Tsiatis, A. A., Davidian, M., Zhang, M., & Laber, E. (2012). Estimating optimal treatment regimes from a classification perspective. *Stat*, *1*, 103–114. <https://doi.org/10.1002/sta.411>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.