



Do machine learning methods lead to similar individualized treatment rules? A comparison study on real data

Florie Bouvier, Etienne Peyrot, Alan Balendran, Corentin Ségalas, Ian Roberts, François Petit, Raphaël Porcher

► To cite this version:

Florie Bouvier, Etienne Peyrot, Alan Balendran, Corentin Ségalas, Ian Roberts, et al.. Do machine learning methods lead to similar individualized treatment rules? A comparison study on real data. *Statistics in Medicine*, 2024, 10.1002/sim.10059 . hal-04503566

HAL Id: hal-04503566

<https://hal.science/hal-04503566>

Submitted on 4 Apr 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Do machine learning methods lead to similar individualized treatment rules? A comparison study on real data.

Florie Bouvier ^{*1}, Etienne Peyrot¹, Alan Balendran¹, Corentin Ségalas², Ian Roberts³, François Petit¹, and Raphaël Porcher^{1,4}

¹Université Paris Cité and Université Sorbonne Paris Nord, Inserm, INRAE, Center for Research in Epidemiology and Statistics (CRESS), F-75004 Paris, France

²Université Bordeaux, Inserm, Bordeaux Population Health Research Center, Bordeaux, France

³Clinical Trials Unit, London School of Hygiene & Tropical Medicine, London, UK

⁴Centre d'Épidémiologie Clinique, Assistance Publique-Hôpitaux de Paris, Hôtel-Dieu, Paris, France

Abstract

Identifying patients who benefit from a treatment is a key aspect of personalized medicine, which allows the development of individualized treatment rules (ITRs). Many machine learning methods have been proposed to create such rules. However, to what extent the methods lead to similar ITRs, i.e., recommending the same treatment for the same individuals is unclear. In this work, we compared 22 of the most common approaches in two randomized control trials. Two classes of methods can be distinguished. The first class of methods relies on predicting individualized treatment effects from which an ITR is derived by recommending the treatment evaluated to the individuals with a predicted benefit. In the second class, methods directly estimate the ITR without estimating individualized treatment effects. For each trial, the performance of ITRs was assessed by various metrics, and the pairwise agreement between all ITRs was also calculated. Results showed that the ITRs obtained via the different methods generally had considerable disagreements regarding the patients to be treated. A better concordance was found among akin methods. Overall, when evaluating the performance of ITRs in a validation sample, all methods produced ITRs with limited performance, suggesting a high potential for optimism. For non-parametric methods, this optimism was likely due to overfitting. The different methods do not

^{*}Correspondence to: Florie Bouvier (florie.brion-bouvier@u-paris.fr)

Hôpital Hôtel-Dieu, 1 place du Parvis de Notre-Dame, 75004 Paris, France

lead to similar ITRs and are therefore not interchangeable. The choice of the method strongly influences for which patients a certain treatment is recommended, drawing some concerns about their practical use.

Keywords: personalized medicine; individualized treatment rule; machine learning; comparison study

1 Introduction

Personalized medicine aims at tailoring a treatment strategy to the individual characteristics of each patient. An essential part of personalized medicine is identifying patients benefiting from a given treatment which allows the construction of individual treatment rules (ITRs). Briefly, ITRs are decision rules that recommend treatment based on patients' characteristics. Of particular interest are optimal treatment rules, which are rules that would lead to the best average outcome in the population if they were followed by all individuals.¹

ITRs can be developed using data from randomized controlled trials (RCTs) or observational data. For instance, Farooq et al. developed the SYNTAX score II to guide decision-making between coronary artery bypass graft surgery (CABG) and percutaneous coronary intervention (PCI) in patients with complex coronary artery disease using data from the SYNTAX trial.² In this paper, we decided to focus on ITRs built from RCTs' data to avoid having to additionally handle confounding factors. However, all the approaches presented here could also be used with observational data.

The PATH statement outlines guidelines for conducting predictive analyses of heterogeneous treatment effects (HTE) in clinical trials. It establishes criteria for predicting HTE and thus developing ITRs, emphasizing the use of a risk modeling approach.^{3,4} However, alternative approaches have also been employed for predicting HTE,⁵ and a myriad of methods whose goal is to construct an ITR has been proposed in the last decade. Nonetheless, their relative performance is not clearly established and, more importantly, it is not clear whether the derived ITRs would lead to recommending the same treatment for the same individuals. This issue is worth studying because if the ITRs are not similar, it is important to know upstream when choosing a method to derive an ITR in real life.

From a conceptual viewpoint, two classes of methods to develop an ITR can be distinguished. The first class relies on deriving individualized treatment effects (ITE) and then an ITR by recommending treatment to those with a predicted benefit. This class can be further divided into two sub-classes: methods estimating the response surfaces and methods directly estimating the ITE via a contrast function. The second class comprehends methods that directly estimate the ITR without explicitly relying on estimating ITEs or a contrast function.

Some comparisons of methods constructing an ITR via the ITE have been performed in the past. In particular, Jacob and Zhang et al. have both studied the performance of meta-learners (T-learner, S-learner, X-learner, DR-learner and R-learner) and causal forests.^{6,7} Jacob found that the methods resulted in differences in terms of ITEs estimates and recommended using multiple methods and comparing their results in practice.⁶ In their paper, Zhang et al. also found that the methods performed differently.⁷ To our knowledge, no comparison has included all the methods we are presenting, par-

ticularly methods that directly construct an ITR without calculating ITEs. Furthermore, none of the previous works implemented metrics to assess the agreements between pairs of methods.

In this study, we aimed to compare a wide range of methods to develop ITRs, both in terms of performance and agreement. We compared 22 different methods and applied those using data from two randomized controlled trials: the International Stroke Trial (IST) and the CRASH-3 trial. The remainder of this paper is organized as follows. We start by introducing the statistical setting in section 2. In section 3, the different methods are presented. In section 4, the 22 methods are compared on two real RCTs. Section 5 concludes with a discussion.

2 Statistical setting

In this section, the potential outcomes framework is introduced. Then, we explain how to construct an individualized treatment rule (ITR). Finally, we enumerate the metrics used to compare the ITRs.

2.1 Causal framework

We follow Rubin’s potential outcomes framework.⁸ We assume access to an independent and identically distributed sample of observations. Let $X \in \mathcal{X} \subset \mathbb{R}^n$ a vector of covariate in the covariate space \mathcal{X} , $A \in \{0, 1\}$ be an indicator variable for the treatment of interest and $Y \in \{0, 1\}$ be a binary outcome. We introduce potential outcomes Y^0 and Y^1 that represent the binary outcomes that would be observed if patients were assigned to either the control or the evaluated treatment respectively. Without loss of generality, we assume that $Y = 1$ is a desirable event. We make the following assumptions:⁹

- Consistency: the observed outcome corresponds to the potential outcome i.e. if a patient received the treatment their observed outcome would be Y^1 and if they received the control, their observed outcome would be Y^0 .
- No interference: the outcome only depends on the treatment applied to the patient, and not on the treatment applied to other patients.
- Unconfoundedness: all characteristics associated with both the treatment assignment and the outcome, should have been measured in the study.
- Positivity: all patients have a non-null probability of receiving either treatment.

In the setting of RCTs, Unconfoundedness and Positivity are met by design.

2.2 Individualized treatment rules

We are interested in constructing individualized treatment rules (ITR) which are decision rules that recommend treatment based on patients’ characteristics.

Those rules are modeled as maps $r : \mathcal{X} \rightarrow \{0, 1\}$. Accordingly, for a given set of covariates

$x \in \mathcal{X}$, $r(x)$ indicates whether or not the treatment should be given to a patient. An optimal rule r^{opt} is obtained when the value $\mathcal{V}(r)$ among all $r \in \mathcal{R}$, with \mathcal{R} being the class of all treatment rules, is maximized:¹

$$r^{opt} = \arg \max_{r \in \mathcal{R}} \mathcal{V}(r),$$

where $\mathcal{V}(r) = E[Y(r)]$ with $Y(r) = Y^1 r(x) + Y^0 [1 - r(x)]$ representing the outcome observed if the rule r was followed.

Constructing an optimal treatment rule can be achieved in two different ways. The first approach involves calculating individual treatment effects (ITE). The ITE τ represents the predicted benefit under one treatment minus the predicted benefit under the other treatment, given a set of patients' characteristics:

$$\tau(x) = E(Y^1 - Y^0 | X = x) = \mu_1(x) - \mu_0(x).$$

An optimal rule is obtained by only giving the evaluated treatment to patients with a positive value of $\tau(x)$ i.e. $r^{opt}(x) = \mathbb{1}_{\{\tau > 0\}}(x)$. In this approach, some methods estimate the ITE by estimating the response surfaces whereas others directly estimate the ITE via a contrast function. The second approach consists of directly developing an optimal rule, without estimating the ITEs, by minimizing a loss function of the value of the rule. The methods to develop ITRs considered in this project are described in section 3.

2.3 Metrics

Several metrics were used to compare the ITRs developed with different methods described in section 3. Using several metrics allows us to have a comprehensive view of the performance of the ITRs. Two classes of metrics can be distinguished: metrics whose aim is to estimate the performance of the rules and metrics whose aim is to compare the level of agreement between two rules.

2.3.1 Performance metrics

First, metrics to assess the quality of a single ITR were used, enabling us to compare the performance of the ITRs.

- The value of a rule: As stated previously, the value $\mathcal{V}(r) = E[Y(r)]$ represents the mean outcome if the ITR was correctly followed. In this project, a desirable binary outcome is considered, thus, ITRs with $\mathcal{V}(r)$ closer to 1 have a better performance.
- The benefit of the rule in terms of assigned treatment among people with a positive and negative score, is assessed with two metrics: B_{pos} and B_{neg} , where B_{pos} represents the average benefit of giving the evaluated treatment among people with a positive score i.e $r(x) = 1$ and B_{neg} represents the average benefit of not giving the evaluated treatment among people with a negative score i.e $r(x) = 0$.¹⁰ The values are between -1 and 1 , with 1 meaning there is a benefit in treating people with a positive score for B_{pos} and a benefit in not treating people with a negative score for B_{neg} .

$$B_{pos} = P(Y = 1 | A = 1, r(x) = 1) - P(Y = 1 | A = 0, r(x) = 1),$$

$$B_{neg} = P(Y = 1 \mid A = 0, r(x) = 0) - P(Y = 1 \mid A = 1, r(x) = 0).$$

- The Population Average Prescription Effect (PAPE): PAPE compares an ITR with a treatment rule that randomly treats the same proportion of patients:¹¹

$$PAPE = E[Y(r) - p_r Y^1 - (1 - p_r) Y^0]$$

where p_r represents the proportion of patients assigned to the evaluated treatment under the ITR r .

The PAPE takes values between -1 and 1 . Here, since higher values of the outcome are desirable, higher values of PAPE indicate a better performance of the ITR. A value of 0 indicates that the ITR does not perform better than treating randomly the same proportion of patients. Negative values mean that the ITR performs worse. An advantage of the PAPE is that it is easy to interpret.

- The c-statistic for benefit: it is the probability that from two randomly chosen matched pairs with unequal observed benefit, the pair with greater observed benefit also has a higher predicted probability where the observed benefit refers to the difference in outcomes between two patients with the same predicted benefit but with different treatment assignments.¹² To create the pairs, a patient in the control group is matched to one in the treatment group with a similar predicted treatment benefit. Higher values of the c-statistic for benefit are better. The c-statistic for benefit quantifies how well the rule discriminates patients benefiting from patients not benefiting from taking a given treatment. The c-statistic for benefit can only be calculated for methods using an ITE or a benefit score to derive an ITR.

The standard errors for each metric were calculated through a Bootstrap procedure involving 1000 samples of the original dataset.

2.3.2 Agreement between two rules

Metrics to see if two ITRs have the same recommendation and agree to allocate the treatment to the same patients were used.

- Matthews correlation coefficient (MCC): Here, the MCC is used to measure the disagreement, in terms of treated patients, between two rules.¹³ The values range between -1 and 1 , where 1 indicates a perfect positive correlation, 0 indicates no correlation and -1 indicates a perfect negative correlation.
- Cohen's kappa coefficient: Cohen's kappa measures the agreement between two rules by considering the number of agreements and disagreements.¹⁴ It can range from -1 to 1 . A value inferior to 0 demonstrates that there is less than chance agreement between the two rules, a value of 0 shows no agreement and a value of 1 means that there is perfect agreement between the rules.

2.4 Multiple Correspondence Analysis

A Multiple Correspondence Analysis (MCA) was conducted to see if the ITRs agreed on the treatment decision in the presence of some specific characteristics. All variables included in the different models were put in the MCA, as well as the treatment allocation recommended by each ITR. Continuous variables were categorized and the choice of the categories was motivated by previous works done using the datasets.^{15,16}

3 Methods to construct individualized treatment rules

This section presents the different methods that were compared. We selected the most common methods which are either simple to implement with the R software¹⁷ or for which a package is available. As mentioned in Section 2.2, two classes of methods to develop an ITR were distinguished: a first class in which the ITR is constructed by first modeling the ITE, and an optimal rule is found by giving the treatment evaluated to the individuals with a positive ITE, and a second class, in which the ITR is directly estimated without calculating individualized treatment effects, and where an optimal rule is found by minimizing the risk of the value of the rule via a loss function. In the first class, two distinct approaches can be used to obtain the ITE: either estimating the expected difference of the potential outcomes between treatments or estimating the ITE directly via a contrast function. The majority of the methods fell under the first category: the meta-learners (T-learner, S-learner, X-learner, DR-learner, and R-learner, both with parametric and non-parametric models), PATH, causal forests, virtual twins, A-learning and the modified covariate method, whereas outcome weighted learning and contrast weighted learning fell under the second class. A classification of methods based on how they construct an optimal treatment rule is given in Figure 1. Conceptually, some methods are related and are therefore referred to as belonging to the same family (e.g. parametric meta-learners, non-parametric meta-learners, A-learning, and the modified covariate method).

3.1 Meta-learners

Meta-learners are methods that use sub-regression problems to estimate the ITE via a base learner. In this project, two base learners were implemented: logistic regression and random forest (RF), the latter is selected for its good performance on tabular data.¹⁸ Since meta-learners can use several base learners, they are flexible and can adapt to different types of data.

Meta-learners with a non-parametric model as the base learner can be prone to overfitting. A solution to this potential overfitting is to use cross-fit.¹⁹ Cross-fit consists of splitting the dataset into several folds. Then, the ITEs are learned on every fold and the results are aggregated to derive an ITR (more details in Supplementary Material S1). The meta-learners were compared with and without cross-fit when using random forests as a base learner. When cross-fit was applied, 5 folds and 30 splits were used, because such a choice has been reported as leading to a good performance.²⁰ However, to our knowledge, there is no standard method or clear guidance on how to perform cross-fit, and other




Figure 1: Classification of the methods.

choices exist.²⁰

In the upcoming segments, we use $\hat{\tau}$ to represent the estimate of τ , and adhere to the same convention for denoting the estimates of other parameters.

3.1.1 S-learner

The S-learner estimates the treatment effect within a single regression model, where the treatment is included as a feature and where interactions between the treatment and relevant covariates are introduced in the parametric setting.²¹ First, use a model to estimate the response function $\mu(x, a)$:

$$\mu(x, a) = \mathbb{E}(Y|X = x, A = a).$$

Then, estimate the individual treatment effect τ :

$$\hat{\tau}(x) = \hat{\mu}(x, 1) - \hat{\mu}(x, 0).$$

3.1.2 T-learner

In the T-learner algorithm, two models are built, one in the treatment group and one in the control group.²¹ These models are used to calculate the response functions:

$$\mu_0(x) = \mathbb{E}(Y|X = x, A = 0),$$

$$\mu_1(x) = \mathbb{E}(Y|X = x, A = 1).$$

The ITE is estimated as the difference between the two predicted risks:

$$\hat{\tau}(x) = \hat{\mu}_1(x) - \hat{\mu}_0(x).$$

3.1.3 X-learner

The X-learner consists of three steps:²¹

1. Estimate the response functions as in the T-learner:

$$\mu_0(x) = \mathbb{E}(Y|X = x, A = 0),$$

$$\mu_1(x) = \mathbb{E}(Y|X = x, A = 1).$$

2. Impute the treatment effects for the individuals in the treated group based on the control-outcome estimator and the treatment effects for the individuals in the control group based on the treatment-outcome estimator and estimate $\hat{\tau}_1(x)$ and $\hat{\tau}_0(x)$:

$$\widetilde{D}^1 = Y^1 - \hat{\mu}_0(X^1),$$

$$\widetilde{D}^0 = \hat{\mu}_1(X^0) - Y^0,$$

$$\hat{\tau}_1(x) = \mathbb{E}(\widetilde{D}^1|X = x),$$

$$\hat{\tau}_0(x) = \mathbb{E}(\widetilde{D}^0|X = x).$$

3. Define the ITE by a weighted average of the two estimates:

$$\hat{\tau}(x) = w(x)\hat{\tau}_0(x) + (1 - w(x))\hat{\tau}_1(x)$$

where $w(x) \in [0, 1]$ is a weighting function. An estimate of the propensity scores can be chosen as the weighting function, but there is no clear theory on how to choose an optimal weighting function. In the setting of RCTs, it is natural to choose $w(x) = \frac{1}{2}$ if the trial had a 1:1 randomization ratio.

The X-learner has been described as advantageous in an unbalanced design or with sparse treatment effects.²¹

3.1.4 DR-learner

The DR-learner is a doubly robust estimator that estimates the ITE in two stages.²² This learner includes double sample splitting to reduce bias.

First, the data $Z_i = (X_i, A_i, Y_i)$ are randomly split into three independent samples D_{1a}, D_{1b}, D_2 . Then, the following two steps are applied:

1. Construct propensity score estimates $\hat{\pi}$ of the propensity scores $\pi(X) = \mathbb{P}(A = 1|X = x)$ using D_{1a} and estimate the response functions $\hat{\mu}_0$ and $\hat{\mu}_1$ using D_{1b} .
2. Construct the pseudo-outcome:

$$\hat{\varphi}(Z) = \frac{A - \hat{\pi}(X)}{\hat{\pi}(X)[1 - \hat{\pi}(X)]} [Y - \hat{\mu}_A(X)] + \hat{\mu}_1(X) - \hat{\mu}_0(X).$$

Then, regressing it on covariates X of D_2 to estimate the ITE:

$$\hat{\tau}(x) = \mathbb{E}[\hat{\varphi}(Z) | X = x].$$

Cross-fitting can be added as an additional third step:

3. Repeat steps 1 and 2 twice. First, D_{1b} and D_2 are used for step 1 and D_{1a} is used for step 2. Then, D_{1a} and D_2 are used for step 1 and D_{1b} is used for step 2. A final estimate of τ is constructed by averaging the three estimates.

In this work, the propensity score was taken equal to $\frac{1}{2}$ since data from RCTs with a 1:1 randomization ratio were used.

3.1.5 R-learner

The R-learner estimates the ITEs in two steps:²³

1. Fit the response function $\hat{\mu}(x)$ and the propensity scores $\hat{\pi}(x)$ with a base learner.
2. Estimate ITEs by minimizing the R-loss, which uses Robinson's decomposition:²⁴

$$L_R(\tau(x)) = \frac{1}{n} \sum_{i=1}^n [(Y_i - \hat{\mu}(X_i)) - (A_i - \hat{\pi}(X_i))\tau(X_i)]^2 + \Lambda_n(\tau(\cdot))$$

where $\Lambda_n(\tau(\cdot))$ is a regularization term on the complexity of $\tau(\cdot)$.

The response function and the propensity scores can be fitted using a cross-validation procedure and the regularization could be done with a penalized regression such as lasso or ridge, for instance, when logistic regression is used as a base learner.²³ When random forests are used, regularization is achieved by tuning the hyperparameters, particularly by limiting the tree's depth and the number of variables used to build the trees.

3.2 PATH approach

PATH is a risk modeling approach that has been recommended by the PATH statement.^{3,4} This method involves three steps:

1. Fit a regression model (a logistic model with a binary outcome) with the relevant variables to derive the linear predictor
2. Build a model that incorporates the linear predictor, the treatment variable, and the interaction between the linear predictor and treatment to estimate the response functions
3. Derive the ITE based on the response functions

This approach has demonstrated strong performance in diverse scenarios according to previous simulation studies.^{25,26}

3.3 Causal forests

The causal forests algorithm is a special case of generalized random forests (GRF), a flexible and general framework to estimate the ITEs.²⁷ Causal forests extend the original random forest algorithm by borrowing ideas from kernel-based methods and the R-learner.²³

In contrast to the standard random forest algorithm in which a prediction for a new observation is obtained by averaging predictions of each tree, here, the trees are used to compute a weighting scheme similar to kernel-based methods. The trees act as weights between training points and any new observations:

$$\alpha_{bi}(x) = \frac{\mathbb{1}_{X_i \in L_b(x)}}{|L_b(x)|}, \quad \alpha_i(x) = \frac{1}{B} \sum_{b=1}^B \alpha_{bi}(x)$$

where X_i corresponds to the covariates of individual i in the training dataset and $L_b(x)$ corresponds to the set of observations in the training set that fall in the same leaf as x for tree b .

Then, the prediction for a new observation is obtained using the adaptive weights by minimizing the R-loss described above.

Another characteristic of causal forests (and more generally of GRF) is the notion of honesty where the training data is split into two parts: one for constructing the tree and the other (the estimation sample) for estimating leaf values for each tree. In doing so, the estimates are less prone to bias and more consistent. The notion of honesty is similar to employing the crossfit in non-parametric meta-learners.

3.4 Virtual twins

The virtual twins method consists in predicting response probabilities for treatment and control twins for all individuals using counterfactual models.²⁸ The difference in the probabilities is then used as the outcome in a classification or regression tree. A subgroup of individuals defined by a region S of the covariate space \mathcal{X} for which the treatment effect τ is better than a prespecified threshold can be then identified. The two steps are described below:

1. Fit a random forest in which the covariates, the treatment indicator, and treatment-covariates interactions are included to estimate the response function μ and the ITE τ as in S-learner.
2. Build a regression or a classification tree to find the covariates X that are strongly associated with τ to define region S . Define τ^* , a binary variable, as the outcome. When $\tau > c$, $\tau^* = 1$ and when $\tau \leq c$, $\tau^* = 0$. Develop an ITR based on the value of τ^* . Individuals for which $\tau^* = 1$ are placed in the estimated region \hat{S} . The evaluated treatment is given to individuals in \hat{S} .

In this work, we built a classification tree and set c equal to 0.

The enhanced treatment effect $Q(S)$ defined as:

$$Q(S) = (P(Y = 1|A = 1, X \in S) - P(Y = 1|A = 0, X \in S)) - (P(Y = 1|A = 1) - P(Y = 1|A = 0))$$

can be estimated by estimating $P(Y = 1|A = 1, X \in \hat{S})$, $P(Y = 1|A = 0, X \in \hat{S})$, $P(Y = 1|A = 1)$ and $P(Y = 1|A = 0)$ using the observed proportions of the data. One of the different approaches that can be used to correct bias is bootstrapping. Bootstrapping measures the bias of $Q(\hat{S})$, which is then used to adjust $Q(\hat{S})$. In their work, Foster et al. compared bootstrapping to other approaches; their conclusion favored bootstrap with 20 samples.²⁸

3.5 A-learning and the modified covariate method

A-learning and the modified covariate method are two methods that focus on treatment-covariates interactions since treatment selection solely depends on the sign of the interactions.^{29,30} Given the covariates and the treatment, the estimated outcome can be written as:

$$E(Y | A, X) = m(X) + A\Delta(X)$$

where $m(X)$, $\Delta(X)$ represent respectively the main effect of X and the treatment effect given X . Only the signs of $\Delta(X)$ matter for treatment selection.

In both methods, a personalized benefit score model f is calculated and its sign, which is consistent with the direction of the treatment effect, is used to construct an ITR. An optimal ITR is found for both methods by minimizing a certain loss function ℓ . Details on the loss functions are given below.

3.5.1 A-learning

In A-learning, the following expected loss function is considered:

$$\ell_A(f) = E(\ell_A(f, x))$$

with

$$\begin{aligned} \ell_A(f, x) = & \pi(x)E[M(Y, (1 - \pi(x))f(x)) | A = 1, X = x] \\ & + (1 - \pi(x))E[M(Y, -\pi(x)f(x)) | A = 0, X = x]. \end{aligned}$$

where $\pi(x)$ represents the propensity scores and M is a positive function, such as the quadratic or cross-entropy (also called logistic loss).

$\ell_A(f, x)$ is then replaced by its empirical version on the observed data:

$$L_A(f) = \frac{1}{n} \sum_{i=1}^n M(Y_i, (A_i - \pi(X_i))f(X_i))$$

$\pi(X_i)$ equals $\frac{1}{2}$ in the context of RCTs with 1:1 randomization.

When M is chosen to be the logistic loss, $L_A(f)$ is expressed as:

$$L_A(f) = -\frac{1}{n} \sum_i Y_i(A_i - \pi(X_i))f(X_i) - \log(1 + \exp((A_i - \pi(X_i))f(X_i))).$$

3.5.2 Modified covariate method

Similarly, the expected loss function $\ell_{MCM}(f) = E(\ell_{MCM}(f, x))$ of the modified covariate method where

$$\begin{aligned}\ell_{MCM}(f, x) = & E[M(Y, f(X)) \mid A = 1, X = x] \\ & + E[M(Y, -f(X)) \mid A = 0, X = x]\end{aligned}$$

Its empirical version is

$$L_{MCM}(f) = \frac{1}{n} \sum_{i=1}^n \frac{M(Y_i, (2A_i - 1)f(X_i))}{(2A_i - 1)\pi(X_i) + 1 - A_i},$$

which boils down to

$$L_{MCM}(f) = -\frac{2}{n} \sum_i Y_i(2A_i - 1)f(X_i) - \log(1 + \exp((2A_i - 1)f(X_i)))$$

when the logistic loss function is used and $\pi(X_i) = \frac{1}{2}$. It is worth mentioning that by substituting the benefit score of A-learning in L_A with double the benefit score of the modified covariate method, we obtain $L_{MCM} = 2L_A$.

3.6 Outcome weighted learning

Outcome weighted learning (OWL) uses a weighted classification framework, in which each patient is weighted based on their outcome, with a hinge loss to estimate an ITR.³¹ An optimal treatment rule is obtained by minimizing the following quantity:

$$L(f) = \frac{1}{n} \sum_{i=1}^n \frac{Y_i}{(2A_i - 1)\pi(X_i) + 1 - A_i} (1 - (2A_i - 1)f(X_i))^+ + \lambda_n \|f\|^2$$

where $x^+ = \max(x, 0)$, $\pi(X_i)$ represents the propensity scores, λ_n is a penalty parameter used to avoid overfitting and $\|f\|$ is a norm for the function $f : \mathcal{X} \rightarrow \{0, 1\}$. When employing the linear kernel, the Euclidean norm of the coefficients, excluding the intercept, is utilized. The parameter λ_n is chosen by performing a cross-validation. OWL is a consistent estimator and has low variability.³¹

3.7 Contrast weighted learning

The idea behind contrast weighted learning (CWL) is to use contrasts of the outcome between pairs of patients to build weights used in a weighted classification algorithm to estimate an ITR.³² A contrast function h is defined for a pair of patients to measure the relative favorability of their outcomes. Several contrast functions exist such as the difference $h(Y_i, Y_j) = Y_i - Y_j$, the log ratio $h(Y_i, Y_j) = \log(Y_i/Y_j)$ or the win indicator $h(Y_i, Y_j) = \text{sgn}(Y_i - Y_j)$, with $\text{sgn}(Y_i - Y_j) = 1$ if $Y_i - Y_j > 0$; $\text{sgn}(Y_i - Y_j) = 0$ if $Y_i - Y_j = 0$ and $\text{sgn}(Y_i - Y_j) = -1$ if $Y_i - Y_j < 0$. In this project, we used the win indicator, considered the most robust contrast function by Guo et al.³² The optimal ITR is found by minimizing the following function:

$$L(f) = \frac{1}{2} E[(\mathbb{1}_{h(Y_i, Y_j)(2A_i - 1)f(X_i) < 0} + \mathbb{1}_{h(Y_i, Y_j)(2A_j - 1)f(X_j) \geq 0}) \frac{|h(Y_i, Y_j)|}{\pi(X_i)\pi(X_j)}]$$

where $h(Y_i, Y_j)$ is the contrast function between patient i and patient j and $\pi(X_j)$ represents the propensity score. Here $h(Y_i, Y_j) = \text{sgn}(Y_i - Y_j)$ and $\pi(X_i) = \pi(X_j) = 1/2$. CWL is a flexible and robust method that only relies on the contrast of outcomes between two patients. However, a correctly specified model is needed to ensure consistency.

3.8 Implementation

All the analyses were performed in R version 4.1.2. Virtual Twins was implemented using the `aVirtualTwins` package.³³ The package `personalized` was used to develop the modified covariate method and A-learning.³⁴ For outcome weighted learning and contrast weighted learning, the package `WeightSVM` was used.³⁵ Causal forests was implemented using the package `grf`.³⁶ More details about the implementation such as the choice of the hyper-parameters can be found in Supplementary Material S2.

4 Comparison of the methods on real data

In this section, the ITRs obtained when applying the methods described above are compared on two multi-center randomized control trials: the International Stroke Trial and the CRASH-3 trial. The train and test datasets were obtained by splitting the data at the center level using 2/3 of the data for training.

4.1 International Stroke Trial

The 22 methods were first compared on the International Stroke Trial (IST).¹⁵ The IST was chosen because an ITR has been developed on this dataset in the past using the T-learner method and found that 74% of patients would benefit from taking aspirin.³⁷ The IST is a multi-center randomized control trial that includes 19,435 patients recruited from 466 centers and examines the impact of administering aspirin, heparin, or both in stroke. For our illustration, we focused on the impact of aspirin on stroke. Nineteen variables were included in the different methods: 16 categorical variables and 3 continuous variables, similar to what Nguyen et al. did.³⁷ The outcome used was death or dependency at 6 months (1=no and 0=yes). The treatment variable was binary (0=no aspirin and 1=aspirin). A description of the covariates and the outcome is reported in Supporting Material S3.

Results of the metrics used to evaluate the ITR produced by each method are given in Table 1. The performance in the train dataset of IST can be found in Supplementary Material S3. Higher values of c-statistic for benefit are better but it is rare to obtain values above 0.6.¹² Here, the c-statistic for benefit was close to 0.5 for all the ITRs, indicating poor discrimination. A reason for the poor discrimination could be the lack of strong heterogeneous treatment effects. This hypothesis was confirmed

by conducting a likelihood ratio test comparing models with and without treatment-covariate interactions were compared. The test showed that the interactions did not add value meaning no significant heterogeneity was found in the IST dataset. The adequacy index was also computed to see how much predictive information was due to the treatment-covariate interactions. For IST, the adequacy index was equal to 0.993 meaning adding the interactions only accounted for 0.7% of the predictive information. The ITRs had a PAPE close to 0 meaning that the ITRs did not perform better than a rule that treated randomly the same proportion of patients. The PAPE values of most meta-learners were even slightly negative indicating that a non-individualized rule performed slightly better than those individualized rules. The values of B_{pos} and B_{neg} were close to 0, meaning there were not many benefits of giving the evaluated treatment to patients with a positive score or not giving the evaluated treatment to patients with a negative score. The proportion of patients for whom aspirin was recommended by the different ITRs ranged from 0.114 to 0.899, with most methods producing an ITR that recommended the evaluated treatment for more than 50% of patients. Methods belonging to the same family had similar proportions. Despite the significant disparity of proportions, the estimated values of the ITRs were similar showing that giving the evaluated treatment to more or fewer patients did not improve the value. For instance, OWL's ITR recommended treating 0.898 of patients, and CWL's ITR recommended treating 0.114 of patients but their rule's values were 0.399 and 0.400 respectively. The mean outcome when no one was treated (0.380) was close to the mean outcome when everyone was treated (0.396), which further implies that the treatment had a limited impact on the outcome on average. The mean outcome under the individualized rule was above the mean outcome when no one was treated for all methods. However, only five methods (Causal forests, A-learning, modified covariate method, OWL, and CWL) had mean outcome under the rule above the mean outcome when everyone was treated, and even for those methods, the mean outcome did not notably surpass the mean outcome if everyone is treated. Generally, the ITRs developed by the different methods did not drastically improve the mean outcome.

Overall, MCC and kappa's coefficient produced similar values (Figure 2). Most methods had considerable disagreements and thus almost no correlation regarding the people treated with the evaluated treatment in their rules which can indicate that the rules did not consider the same characteristics for the allocation of the treatment. A better concordance was found among methods of the same family. For instance, the ITRs developed with the parametric meta-learners agreed to treat similar patients and had MCC and Cohen's kappa values ranging from 0.77 to 1. Similarly, non-parametric meta-learners had a positive moderate to high correlation with each other and with their crossfitted counterparts. However, they had less correlation with the parametric ones. A-learning and the modified covariate method generated the same ITR and therefore had coefficients of 1. The ITRs obtained with the different methods generally did not recommend the evaluated treatment to the same patients, which draws some concerns for their usability in practice.

A majority of characteristics were located near the origin and were not associated with the treatment allocation of the different ITRs (Figure 3). Virtual twins' ITR recommended not treating patients in a

Table 1: Results of the metrics for each method applied to the IST dataset.

	p_r	$\mathcal{V}(r)$ (SE)	$E(Y^0)$ (SE)	$E(Y^1)$ (SE)	B_{pos} (SE)	B_{neg} (SE)	PAPE (SE)	c for benefit (95% CI)
SL	0.567	0.388 (0.009)	0.380 (0.009)	0.396 (0.009)	0.014 (0.017)	-0.019 (0.018)	-0.002 (0.006)	0.495 (0.476;0.514)
TL	0.567	0.388 (0.009)	0.380 (0.009)	0.396 (0.009)	0.014 (0.017)	-0.019 (0.019)	-0.002 (0.006)	0.495 (0.476;0.514)
XL	0.622	0.382 (0.009)	0.380 (0.009)	0.396 (0.009)	0.004 (0.016)	-0.034 (0.020)	-0.008 (0.006)	0.495 (0.477;0.513)
DRL	0.622	0.382 (0.008)	0.380 (0.009)	0.396 (0.009)	0.003 (0.016)	-0.035 (0.020)	-0.008 (0.006)	0.495 (0.477;0.514)
RL	0.626	0.388 (0.009)	0.380 (0.009)	0.396 (0.009)	0.013 (0.016)	-0.022 (0.020)	-0.003 (0.006)	0.501 (0.483;0.519)
SL RF	0.526	0.395 (0.009)	0.380 (0.009)	0.396 (0.009)	0.030 (0.017)	-0.004 (0.017)	0.006 (0.006)	0.498 (0.480;0.516)
TL RF	0.515	0.390 (0.009)	0.380 (0.009)	0.396 (0.009)	0.021 (0.018)	-0.013 (0.018)	0.002 (0.006)	0.498 (0.480;0.517)
XL RF	0.543	0.393 (0.009)	0.380 (0.009)	0.396 (0.009)	0.025 (0.017)	-0.008 (0.017)	0.004 (0.006)	0.504 (0.486;0.523)
DRL RF	0.529	0.390 (0.009)	0.380 (0.009)	0.396 (0.009)	0.020 (0.017)	-0.015 (0.018)	0.001 (0.006)	0.500 (0.482;0.518)
RL RF	0.528	0.387 (0.009)	0.380 (0.009)	0.396 (0.009)	0.018 (0.018)	-0.019 (0.018)	-0.001 (0.006)	0.501 (0.482;0.519)
SL CF	0.598	0.389 (0.009)	0.380 (0.009)	0.396 (0.009)	0.018 (0.017)	-0.019 (0.019)	0.000 (0.006)	0.497 (0.478;0.515)
TL CF	0.555	0.385 (0.009)	0.380 (0.009)	0.396 (0.009)	0.009 (0.017)	-0.025 (0.019)	-0.004 (0.006)	0.498 (0.480;0.516)
XL CF	0.599	0.386 (0.009)	0.380 (0.009)	0.396 (0.009)	0.011 (0.017)	-0.024 (0.018)	-0.004 (0.006)	0.498 (0.480;0.516)
DRL CF	0.552	0.386 (0.009)	0.380 (0.009)	0.396 (0.009)	0.011 (0.016)	-0.023 (0.019)	-0.003 (0.006)	0.499 (0.481;0.517)
RL CF	0.548	0.389 (0.009)	0.380 (0.009)	0.396 (0.009)	0.018 (0.018)	-0.016 (0.017)	0.000 (0.006)	0.501 (0.483;0.519)
PATH	0.723	0.394 (0.009)	0.380 (0.009)	0.396 (0.009)	0.020 (0.015)	-0.003 (0.014)	0.002 (0.005)	0.511 (0.493;0.529)
Causal Forests	0.899	0.397 (0.008)	0.380 (0.009)	0.396 (0.009)	0.019 (0.013)	0.010 (0.032)	0.002 (0.003)	0.504 (0.486;0.522)
VT	0.758	0.393 (0.009)	0.380 (0.009)	0.396 (0.009)	0.017 (0.015)	-0.009 (0.019)	0.000 (0.005)	—
AL	0.501	0.419 (0.008)	0.380 (0.009)	0.396 (0.009)	0.089 (0.018)	0.059 (0.018)	0.031 (0.005)	0.566 (0.548;0.583)
MCM	0.501	0.419 (0.008)	0.380 (0.009)	0.396 (0.009)	0.089 (0.017)	0.059 (0.017)	0.031 (0.005)	0.566 (0.548;0.583)
OWL	0.898	0.399 (0.009)	0.380 (0.009)	0.396 (0.009)	0.022 (0.013)	0.000 (0.017)	0.004 (0.003)	—
CWL	0.114	0.400 (0.009)	0.380 (0.009)	0.396 (0.009)	0.075 (0.034)	0.007 (0.012)	0.018 (0.004)	—

p_r refers to the proportion of patients for which the rule recommends treatment. $E(Y^0)$ and $E(Y^1)$ refer respectively to the mean outcome when no one is treated and the mean outcome when everyone is treated. SL: S-learner, TL: T-learner, XL: X-learner, DRL: DR-learner, RL: R-learner, RF: random forests, CF: cross-fitted, VT: virtual twins, MCM: modified covariate method, AL: A-learning, OWL: outcome weighted learning and CWL: contrast weighted learning.




Figure 2: Heatmap representing the MCC and Cohen’s Kappa for each combination of two ITRs using the International Stroke Trial.

drowsy state and patients with a TACS stroke and recommended treating fully alert patients, younger patients, and patients with no deficit or disorder. Non-parametric meta-learners produced ITRs that recommended treatment for patients from South America and no treatment for patients from the Middle East, South Asia, and Oceania. CWL’s ITR recommended treating patients with a LACS stroke or another type of stroke while OWL’s ITR recommended not giving aspirin to unconscious patients. The MCA was concordant with what was found in Figure 2 and reflected well the disagreement in terms of treatment allocation between the ITRs.

4.2 CRASH-3

In a second stage, we compared the methods on the CRASH-3 dataset.¹⁶ Some heterogeneity in early treatment administration has been found in the CRASH-3 trial, therefore we thought it would be interesting to develop ITRs on this data.¹⁶ CRASH-3 is a multi-center randomized control trial consisting of 9,072 patients from 175 hospitals over 29 countries. The aim of this trial was to examine the effects of tranexamic acid (TXA) in patients with acute traumatic brain injury. This paper used head injury death as the outcome (1=no and 0=yes). Six covariates were included in the methods: 2 categorical variables and 4 continuous variables. A binary treatment variable (0=Placebo and 1=TXA) was used. A description of the covariates and outcome is given in Supporting Material S3.

Table 2 shows the values of the metrics obtained with each method. The results of the training set are




Figure 3: Multiple Correspondence Analysis on the International Stroke Trial showing all levels of each variable and the treatment recommendation of the individualized treatment rules.

given in Supplementary Material S3. Recall that higher values of c-statistic for benefit are better but it is rare to obtain values above 0.6.¹² Here, c-statistic for benefit values were all around or under 0.5, indicating poor discrimination except for A-learning and the modified covariate method's ITRs which had higher values (0.721 and 0.720 respectively). Excluding A-learning and the modified covariate method's ITRs, the ITRs were not able to differentiate patients benefiting from taking the evaluated treatment from patients not benefiting. The PAPE values were close to 0 meaning that the ITRs did not perform better than a rule which randomly treated the same proportion of patients. There was a mix of positive and negative values but they all remained close to 0. A-learning and the modified covariate method's ITRs had B_{pos} and B_{neg} values around 0.2, showing some benefits of giving the evaluated treatment to patients with a positive score and not giving the evaluated treatment to patients with a negative score, which was not the case for the other ITRs who had values near 0. A-learning and the modified covariate method outperformed other approaches. This better performance is attributed to the treatment rules they developed, predominantly recommending treatment for patients with moderate Glasgow coma scores and reactive pupils. These patients' profiles align with findings from the CRASH-3 study. The proportions of people for which the treatment was recommended went from 0 to 1 with a majority of methods recommending to give the evaluated treatment to over 60% of patients. Note that CWL's ITR chose to give the evaluated treatment to no one whereas OWL's ITR chose to give the evaluated treatment to everyone. The rules' mean outcomes were almost identical and practically all above 0.8, although the proportion of treated patients differed for each method leading us to conclude that there was a negligible treatment effect. This can be emphasized by look-




Figure 4: Heatmap representing the MCC and Cohen’s Kappa for each combination of two ITRs using the CRASH-3 trial.

ing at the mean outcome when no one is treated and the mean outcome when all the individuals are treated. The mean outcome when no one was treated (0.802) was close to the mean outcome when everyone was treated (0.819). Comparing these two mean outcomes to the mean outcome under the rules, we found that all the methods, except the crossfitted S-learner, had a mean outcome higher than the mean outcome when no one was treated, but only three of them (PATH, A-learning, and the modified covariate method) had a better mean outcome than the mean outcome when everyone was treated.

MCC and Cohen’s Kappa coefficient were concordant and gave coefficients of similar magnitude (Figure 4). When one of the ITRs recommended treating everyone or no one with the evaluated treatment, it did not make sense to calculate the MCC and Cohen’s Kappa coefficient, therefore we put a dashed line in those cases. Parametric meta-learners had a strong concordance with each other with high coefficients. The same thing was observed for non-parametric meta-learners whether crossfit was applied or not, as well as for A-learning and the Modified covariate method. As for the IST, a strong concordance is only found between methods belonging to the same family (e.g. parametric meta-learners, non-parametric meta-learners, A-learner and the modified covariate method). Otherwise, the correlation between the ITRs was moderate and most of the time low. The ITRs did not recommend the evaluated treatment to the same patients. The choice of the method had a big impact on the treatment allocation, meaning that in practice two different methods could lead to completely different rules.

Table 2: Results of the metrics for each method applied to the CRASH-3 dataset.

	p_r	$\mathcal{V}(r)$ (SE)	$E(Y^0)$ (SE)	$E(Y^1)$ (SE)	B_{pos} (SE)	B_{neg} (SE)	PAPE (SE)	c for benefit (95% CI)
SL	0.798	0.818 (0.009)	0.802 (0.010)	0.819 (0.010)	0.019 (0.013)	0.007 (0.040)	0.002 (0.007)	0.485 (0.448;0.522)
TL	0.798	0.818 (0.009)	0.802 (0.010)	0.819 (0.010)	0.019 (0.013)	0.007 (0.037)	0.002 (0.007)	0.485 (0.448;0.522)
XL	0.711	0.812 (0.010)	0.802 (0.010)	0.819 (0.010)	0.013 (0.014)	-0.021 (0.031)	-0.002 (0.007)	0.480 (0.447;0.514)
DRL	0.711	0.812 (0.010)	0.802 (0.010)	0.819 (0.010)	0.013 (0.015)	-0.021 (0.030)	-0.002 (0.007)	0.481 (0.448;0.513)
RL	0.711	0.812 (0.010)	0.802 (0.010)	0.819 (0.010)	0.013 (0.014)	-0.021 (0.030)	-0.002 (0.007)	0.481 (0.448;0.513)
SL RF	0.496	0.804 (0.010)	0.802 (0.010)	0.819 (0.010)	0.003 (0.019)	-0.030 (0.020)	-0.006 (0.007)	0.489 (0.454;0.524)
TL RF	0.514	0.805 (0.010)	0.802 (0.010)	0.819 (0.010)	0.007 (0.018)	-0.022 (0.021)	-0.002 (0.007)	0.493 (0.458;0.528)
XL RF	0.579	0.809 (0.010)	0.802 (0.010)	0.819 (0.010)	0.012 (0.017)	-0.022 (0.021)	-0.002 (0.007)	0.490 (0.457;0.524)
DRL RF	0.553	0.813 (0.009)	0.802 (0.010)	0.819 (0.010)	0.019 (0.016)	-0.015 (0.022)	0.001 (0.007)	0.518 (0.484;0.552)
RL RF	0.531	0.812 (0.010)	0.802 (0.010)	0.819 (0.010)	0.018 (0.018)	-0.016 (0.021)	0.001 (0.007)	0.512 (0.479;0.545)
SL CF	0.641	0.798 (0.010)	0.802 (0.010)	0.819 (0.010)	-0.008 (0.015)	-0.056 (0.027)	0.001 (0.007)	0.469 (0.435;0.504)
TL CF	0.605	0.803 (0.010)	0.802 (0.010)	0.819 (0.010)	0.000 (0.015)	-0.041 (0.025)	-0.009 (0.007)	0.476 (0.443;0.510)
XL CF	0.666	0.804 (0.010)	0.802 (0.010)	0.819 (0.010)	0.003 (0.014)	-0.051 (0.028)	-0.010 (0.008)	0.473 (0.440;0.507)
DRL CF	0.593	0.803 (0.010)	0.802 (0.010)	0.819 (0.010)	0.001 (0.017)	-0.039 (0.024)	-0.010 (0.007)	0.485 (0.451;0.518)
RL CF	0.592	0.802 (0.010)	0.802 (0.010)	0.819 (0.010)	-0.002 (0.017)	-0.042 (0.023)	-0.011 (0.007)	0.486 (0.454;0.519)
PATH	0.979	0.820 (0.009)	0.802 (0.010)	0.819 (0.010)	0.019 (0.013)	-0.008 (0.023)	0.001 (0.004)	0.494 (0.462;0.526)
Causal Forests	0.905	0.817 (0.010)	0.802 (0.010)	0.819 (0.010)	0.016 (0.013)	0.003 (0.058)	-0.001 (0.006)	0.471 (0.437;0.506)
VT	0.897	0.819 (0.010)	0.802 (0.010)	0.819 (0.010)	0.019 (0.013)	0.018 (0.050)	0.002 (0.007)	—
AL	0.503	0.859 (0.007)	0.802 (0.010)	0.819 (0.010)	0.236 (0.027)	0.212 (0.027)	0.049 (0.004)	0.721 (0.699;0.743)
MCM	0.503	0.859 (0.007)	0.802 (0.010)	0.819 (0.010)	0.236 (0.027)	0.212 (0.028)	0.049 (0.004)	0.720 (0.698;0.743)
OWL	1	0.819 (0.010)	0.802 (0.010)	0.819 (0.010)	0.017 (0.014)	—	0.000 (0)	—
CWL	0	0.802 (0.010)	0.802 (0.010)	0.819 (0.010)	—	-0.017 (0.013)	0.000 (0)	—

p_r refers to the proportion of patients for which the rule recommends treatment. $E(Y^0)$ and $E(Y^1)$ refer respectively to the mean outcome when no one is treated and the mean outcome when everyone is treated. SL: S-learner, TL: T-learner, XL: X-learner, DRL: DR-learner, RL: R-learner, RF: random forests, CF: cross-fitted, VT: virtual twins, MCM: modified covariate method, AL: A-learning, OWL: outcome weighted learning and CWL: contrast weighted learning.




Figure 5: Multiple Correspondence Analysis of the CRASH-3 trial representing the variables' levels and the treatment recommendation of every individualized treatment rule.

The ITRs developed with the parametric meta-learners, virtual twins and causal forests recommended not to treat patients with a low Glasgow Coma Scale score or/and patients with none or only pupil that reacted, whereas they recommended the treatment to patients with a moderate to high Glasgow Coma Scale score, patients who were female and patients with moderate systolic blood pressure (Figure 5). The non-parametric meta-learners' ITRs recommended treating patients younger patients with relatively high blood pressure and not treating patients with low blood pressure. The MCA reflected well the agreement results that were found in Figure 4. Akin methods' ITRs agreed on the treatment allocation but overall the ITRs did not take into account the same characteristics for the treatment decision.

5 Discussion

This paper compared different methods used to construct individualized treatment rules using data from two RCTs: the International Stroke Trial and the CRASH-3 trial. We considered 22 methods belonging to two different classes. The first class included methods that predicted the ITE to derive an ITR: meta-learners (T-learner, S-learner, X-learner, DR-learner and R-learner, both with logistic regression or random forests as a base learner with and without cross-fit), PATH, causal forests, virtual twins, A-learning and modified covariate method. The second class covered methods that directly estimated the ITR without explicitly estimating ITEs: outcome-weighted learning and contrast weighted learning. For each trial, the performance of the ITRs was assessed with various metrics.

The pairwise agreement between ITRs was also evaluated.

Results showed that the ITRs obtained by the different methods generally had considerable disagreements regarding the individuals to be treated with the evaluated treatment for both trials. The proportions of patients for whom the evaluated treatment was recommended by the rules were very different depending on which method was employed to build the ITR and the Cohen's kappa and Matthews correlation coefficients were low. A better concordance was found among methods of the same family (e.g. among all meta-learners with parametric models, or all meta-learners with non-parametric models and cross-fitting). Overall, when evaluating the performance of ITRs in a hold-out validation sample (33% of the original sample selected at random), results showed that all ITRs had limited performance, whatever the performance in the training set, which suggests a high potential of optimism for the algorithms.

The limited performance results might be due to the distribution of treatment effects and the level of heterogeneity. Although some heterogeneity of treatment effects was found in the trials used in this work, especially in the CRASH-3 trial,¹⁶ the level of heterogeneity might not be sufficient to develop a beneficial individualized treatment rule. This result was reinforced by performing likelihood ratio tests and calculating adequacy indexes. For both trials, the likelihood ratio tests led to the conclusion that there was no evidence of significant heterogeneity and the values of the adequacy indexes showed that the treatment-covariate interactions only accounted for a low percentage of the predictive information. Another explanation for the limited performance might be the sample size. Even if the methods were compared on two large RCTs, perhaps more data is needed to obtain a better performance. A solution might be using individual participant data meta-analyses (IPD-MA) since they include a larger number of participants. However, one should consider the heterogeneity that may arise between the studies included in the meta-analysis. Different methods to tackle the heterogeneity in IPD-MA have been proposed and compared in previous works.^{38–40}

In a previous work, Rekkas et al.²⁵ demonstrated via a simulation study that "complex" methods, which are more flexible, require large sample sizes to perform well and that, when one has access to moderate sample sizes, simpler risk modeling methods recommended by the PATH statement³ should be preferred to obtain a good performance. Using more parsimonious models with fewer covariates, like what has been done for the SYNTAX II score, might also lead to more robust ITRs with better agreements.⁴¹ Investigating for which distribution of treatment effects, a model can have good discrimination, and thus be able to develop a beneficial ITR, as well as the requirements in effective sample size to allow reliable development of ITRs, is worth studying.

Some comparisons of methods used to construct ITRs have been conducted in the past,^{6,7} but to our knowledge, no study has investigated the agreements in terms of the treatment decision with all the methods presented in this project. Both Jacob and Zhang et al. have found that the methods had different performances.^{6,7} These results were concordant with ours.

Although we compared many methods in this work, we did not include every existing method. Indeed, we decided to focus on methods that are commonly used and that are easily computed or for which an R package was available. We also focused on real data, and a simulation study should be conducted to better delineate the parameters associated with a better performance of the methods. A recent sim-

ulation study showed that the sample size and the shape of the distribution of treatments impacted the performance of the methods, particularly the performance of "complex" methods.²⁵ However, we considered that the illustration on two large RCTs was necessary to study the agreement between the different ITRs in real settings, simulations being often over-simplified. Using real data also allows for tailoring each method, in the sense that each model does not necessarily need to have the same variables. Furthermore, in this paper, we decided to compare the ITRs' decisions on randomized controlled trials. Constructing ITRs can also be done with observational data. Observational databases have the potential to include much more participants, and more diverse participants, than trials, and thus might have both more heterogeneity and larger sample sizes, and could be a better source of data to develop ITRs in practice. Although an effort was dedicated to method optimization, specifically optimizing hyperparameters for tree-based methods through cross-validation to maximize accuracy, it is plausible that further optimization might have led to improved method performance.

In conclusion, the significant disagreements that the methods had regarding the treatment allocation suggest that the methods are not interchangeable. Therefore, the chosen method greatly influences the patients for which the evaluated treatment is recommended. It draws some concerns about their practical use. Some ITRs have been developed in the past using one method with similar RCTs.³⁷ Using multiple methods and comparing the obtained ITRs, as suggested by Jacob, might be a solution when one wants to develop an ITR in practice.⁶ However, in most cases, more simple approaches such as the risk modeling method advocated in PATH,³ or carefully adding specific interactions between prespecified treatment-effect modifiers and treatment in the model, as done in the revised SYNTAX score II⁴¹ may be a better strategy than currently available ITR algorithms which may be misleading by overfitting the heterogeneity of treatment effects. Also, methods that allow evaluating the model calibration for benefit may be favored. Evaluating a priori the probability of identifying a beneficial ITR, as suggested by Cain et al.,⁴² might also be taken under consideration.

Declarations

Competing interests

The authors declare that they have no competing interests.

Funding

FB and RP acknowledge support by the French Agence Nationale de la Recherche as part of the "Investissements d'avenir" program, reference ANR-19-P3IA-0001 (PRAIRIE 3IA Institute). This work was partially funded by the Agence Nationale de la Recherche, under grant agreement no. ANR-18-CE36-0010-01.

Authors' contributions

Study concept and design: FB, CS, FP and RP. Analysis and interpretation of data: All authors. Drafting of the manuscript: FB, FP, and RP. Critical revision of the manuscript for important intellectual content: All authors.

Acknowledgements

The authors wish to thank the investigators of the International Stroke Trial and the CRASH-3 trial for providing the datasets compared in this work. The authors would also like to thank Viet-Thi Tran and François Grolleau for contributing to the comparison using the CRASH-3 dataset. Finally, We thank the two anonymous referees for their comments that greatly improved the paper.

Data availability statement

The data that support the findings of this study are openly available in IST dataset at <https://datashare.ed.ac.uk/handle/10283/124> and in CRASH-3 at <https://freebird.lshtm.ac.uk/index.php/available-trials/>

References

1. Zhang B, Tsiatis AA, Laber EB, Davidian M. A robust method for estimating optimal treatment regimes. *Biometrics*. 2012;68(4):1010–1018.
2. Farooq V, Klaveren D, Steyerberg EW, *et al.* Anatomical and clinical characteristics to guide decision making between coronary artery bypass surgery and percutaneous coronary intervention for individual patients: development and validation of SYNTAX score II. *Lancet*. 2013;381(9867):639–650.
3. Kent DM, Steyerberg E, Klaveren D. Personalized evidence based medicine: predictive approaches to heterogeneous treatment effects. *BMJ*. 2018;363.
4. Kent DM, Paulus JK, Klaveren D, *et al.* The Predictive Approaches to Treatment effect Heterogeneity (PATH) Statement. *Annals of Internal Medicine*. 2020;172(1):35–45.
5. Rekkas A, Paulus JK, Raman G, *et al.* Predictive approaches to heterogeneous treatment effects: a scoping review. *BMC Medical Research Methodology*. 2020;20(1):264.
6. Jacob D. CATE meets ML – The Conditional Average Treatment Effect and Machine Learning. 2021.
7. Zhang W, Li J, Liu L. A Unified Survey of Treatment Effect Heterogeneity Modelling and Uplift Modelling. *ACM Computing Surveys*. 2021;54(8).
8. Rubin DB. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*. 1974;6:688–701.
9. Goetghebeur E, Cessie S, De Stavola B, Moodie EE, Waernbaum I, STRATOS initiative . Formulating causal questions and principled statistical answers. *Statistics in Medicine*. 2020;39(30):4922-4948.
10. Janes H, Brown MD, Pepe M, Huang Y. Statistical Methods for Evaluating and Comparing Biomarkers for Patient Treatment Selection. *UW Biostatistics Working Paper Series*. 2013.
11. Imai K, Li ML. Experimental Evaluation of Individualized Treatment Rules. *Journal of the American Statistical Association*. 2021;118(541):242-256.
12. Klaveren D, Steyerberg EW, Serruys PW, Kent DM. The proposed 'concordance-statistic for benefit' provided a useful metric when modeling heterogeneous treatment effects. *Journal of Clinical Epidemiology*. 2018;94:59–68.
13. Baldi P, Brunak S, Chauvin Y, Andersen CAF, Nielsen H. Assessing the accuracy of prediction algorithms for classification: an overview. *Bioinformatics*. 2000;16(5):412–424.
14. Cohen J. A Coefficient of Agreement for Nominal Scales. *Educational and Psychological Measurement*. 1960;20(1):37–46.

15. the International Stroke Trial Collaborative Group , Sandercock PA, Niewada M, Członkowska A. The International Stroke Trial database. *Trials*. 2011;12(1):101.
16. The CRASH-3 trial collaborators . Effects of tranexamic acid on death, disability, vascular occlusive events and other morbidities in patients with acute traumatic brain injury (CRASH-3): a randomised, placebo-controlled trial. *The Lancet*. 2019;394(10210):1713–1723.
17. R Core Team . *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing Vienna, Austria 2022.
18. Grinsztajn L, Oyallon E, Varoquaux G. Why do tree-based models still outperform deep learning on tabular data?. 2022.
19. Chernozhukov V, Chetverikov D, Demirer M, Duflo E, Hansen C, Newey W. Double/Debiased/Neyman Machine Learning of Treatment Effects. 2017.
20. Jacob D. Cross-Fitting and Averaging for Machine Learning Estimation of Heterogeneous Treatment Effects. 2020.
21. Künnel SR, Sekhon JS, Bickel PJ, Yu B. Metalearners for estimating heterogeneous treatment effects using machine learning. *Proceedings of the National Academy of Sciences U S A*. 2019;116(10):4156–4165.
22. Kennedy EH. Optimal doubly robust estimation of heterogeneous causal effects. 2020.
23. Nie X, Wager S. Quasi-oracle estimation of heterogeneous treatment effects. *Biometrika*. 2020;108(2):299-319.
24. Robinson P. Root- N-Consistent Semiparametric Regression. *Econometrica*. 1988;56(4):931-54.
25. Rekkas A, Rijnbeek PR, Kent DM, Steyerberg EW, Van Klaveren D. Estimating individualized treatment effects from randomized controlled trials: a simulation study to compare risk-based approaches. *BMC Medical Research Methodology*. 2023;23(1):74.
26. Van Klaveren D, Balan TA, Steyerberg EW, Kent DM. Models with interactions overestimated heterogeneity of treatment effects and were prone to treatment mistargeting. *Journal of Clinical Epidemiology*. 2019;114:72–83.
27. Athey S, Tibshirani J, Wager S. Generalized Random Forests. 2018.
28. Foster J, Taylor J, Ruberg S. Subgroup identification from randomized clinical trial data. *Statistics in Medicine*. 2011;30:2867–2880.
29. Chen S, Tian L, Cai T, Yu M. A general statistical framework for subgroup identification and comparative treatment scoring. *Biometrics*. 2017;73(4):1199–1209.
30. Tian L, Zhao L, Wei L. Predicting the restricted mean event time with the subject’s baseline covariates in survival analysis. *Biostatistics*. 2014;15:222–233.

31. Zhao Y, Zeng D, Rush AJ, Kosorok MR. Estimating Individualized Treatment Rules Using Outcome Weighted Learning. *Journal of the American Statistical Association*. 2012;107(449):1106–1118.
32. Guo X, Ni A. Contrast weighted learning for robust optimal treatment rule estimation. *Statistics in Medicine*. 2022:sim.9574.
33. Vieille F, Foster J. *aVirtualTwins: Adaptation of Virtual Twins Method from Jared Foster* 2018. R package version 1.0.1.
34. Huling JD, Yu M. Subgroup Identification Using the personalized Package. *Journal of Statistical Software*. 2021;98(5):1–60.
35. Xu T, Chang CC, Lin CC, *et al*. *WeightSVM: Subject Weighted Support Vector Machines* 2022. R package version 1.7-11.
36. Tibshirani J, Athey S, Sverdrup E, Wager S. *grf: Generalized Random Forests* 2022. R package version 2.2.1.
37. Nguyen TL, Collins GS, Landais P, Le Manach Y. Counterfactual clinical prediction models could help to infer individualized treatment effects in randomized controlled trials-An illustration with the International Stroke Trial. *Journal of Clinical Epidemiology*. 2020;125:47–56.
38. Brion Bouvier F, Chaimani A, Gueyffier F, Grenet G, Porcher R. Estimating individualized treatment effects using individual participant data meta-analysis. working paper or preprint 2022.
39. Steyerberg EW, Nieboer D, Debray TPA, Houwelingen HC. Assessment of heterogeneity in an individual participant data meta-analysis of prediction models: An overview and illustration. *Statistics in Medicine*. 2019;38(22):4290–4309.
40. Debray TP, Moons KG, Ahmed I, Koffijberg H, Riley RD. A framework for developing, implementing, and evaluating clinical prediction models in an individual participant data meta-analysis. *Statistics in Medicine*. 2013;32(18):3158–3180.
41. Takahashi K, Serruys PW, Fuster V, *et al*. Redevelopment and validation of the SYNTAX score II to individualise decision making between percutaneous and surgical revascularisation in patients with complex coronary artery disease: secondary analysis of the multicentre randomised controlled SYNTAXES trial with external cohort validation. *The Lancet*. 2020;396(10260):1399–1412.
42. Cain CH, Murray TA, Rudser KD, *et al*. Design considerations and analytical framework for reliably identifying a beneficial individualized treatment rule. *Contemporary Clinical Trials*. 2022;123:106951.