Published in final edited form as:

J Comput Graph Stat. 2018 ; 27(4): 849–860. doi:10.1080/10618600.2018.1451337.

An Algorithm for Generating Individualized Treatment Decision Trees and Random Forests

Kevin Doubleday,

Department of Epidemiology and Biostatistics, University of Arizona

Hua Zhou,

Department of Biostatistics, University of California, Los Angeles

Haoda Fu,

Eli Lilly and Company

Jin Zhou*

Department of Epidemiology and Biostatistics, University of Arizona

Abstract

With new treatments and novel technology available, precision medicine has become a key topic in the new era of healthcare. Traditional statistical methods for precision medicine focus on subgroup discovery through identifying interactions between a few markers and treatment regimes. However, given the large scale and high dimensionality of modern data sets, it is difficult to detect the interactions between treatment and high dimensional covariates. Recently, novel approaches have emerged that seek to directly estimate individualized treatment rules (ITR) via maximizing the expected clinical reward by using, for example, support vector machines (SVM) or decision trees. The latter enjoys great popularity in clinical practice due to its interpretability. In this paper, we propose a new reward function and a novel decision tree algorithm to directly maximize rewards. We further improve a single tree decision rule by an ensemble decision tree algorithm, ITR random forests. Our final decision rule is an average over single decision trees and it is a soft probability rather than a hard choice. Depending on how strong the treatment recommendation is, physicians can make decisions based on our model along with their own judgment and experience. Performance of ITR forest and tree methods is assessed through simulations along with applications to a randomized controlled trial (RCT) of 1385 patients with diabetes and an EMR cohort of 5177 patients with diabetes. ITR forest and tree methods are implemented using statistical software R (<https://github.com/kdoub5ha/ITR.Forest>).

Keywords

Optimization; Precision Medicine; Recursive Partitioning; Subgroup Identification; Value Function; Variable Importance

*Corresponding Author.

SUPPLEMENTARY MATERIAL

Supplementary materials provide details of the algorithms in section 2 and additional graphics from section 3.3. Code used for these analyses is available as an R package at <https://github.com/kdoub5ha/ITR.Forest>

1 Introduction

Many diseases, such as Type 2 diabetes (T2D), have a complex and multifactorial pathophysiology. Treatment of T2D typically begins with lifestyle modification and metformin, a medication that lowers blood glucose by reducing glucose production in the liver and enhancing muscle glucose uptake (Arakaki et al., 2014). When lifestyle modification and metformin are insufficient to control blood glucose, additional medications are necessary. There are several options for second- and third-line therapies for T2D, including sulfonylureas, thiazolidinediones, dipeptidyl peptidase-4 (DPP-4) inhibitors, sodium-glucose co-transporter 2 (SGLT2) inhibitors, glucagon-like peptide-1 (GLP-1) receptor agonists, and insulin. Different therapeutic drug classes have different mechanisms in treating T2D, resulting in some advantages and/or disadvantages, limitations, and adverse effects. Clinical guidelines provide less clarity regarding optimal second- and third-line therapies (Diamant et al., 2014; Forst et al., 2014). The unclear advantages of a combination of drug regimens and the increased potential for adverse effects make glucose lowering therapy increasingly complex (Bergenstal et al., 2010; Nyenwe et al., 2011). So far, there is no consensus for individualized treatment guidance on the selection of these treatments, especially when taking into account treatment heterogeneity effects. There are increasing efforts to develop individualized treatment rules (ITR) in the new era of precision medicine (Hayes et al., 2007; Cummings et al., 2010).

Traditional methods of personalized treatment suggestions are based on ad hoc subgroup analyses or searching for covariate-treatment interactions. These methods, however, may suffer from a lack of efficiency or validity due to the curse of dimensionality and multiple comparisons. Recent methodologies for developing ITR generally fall into three categories. The first approach focuses on developing novel algorithms of covariate-treatment interaction detection. For example, Su et al. (2009) and Lipkovich et al. (2011) developed the interaction tree methods by building splitting rules based on covariate-treatment interaction tests. Su et al. (2012) then extended their previous method to observational studies. The second category is two-step methods (Cai et al., 2011; Zhao et al., 2013; Foster et al., 2011; Faries et al., 2013). The first step is to estimate the differential treatment effect of each individual patient measured by a score function. These scores are then used as responses to establish a relationship with the covariates as the second step. The third class of methods is based on value functions which obtain optimal ITR by maximizing the value function (Qian and Murphy, 2011; Zhao et al., 2012; Zhang et al., 2012a). New methodologies have greatly extended our ability to explore solutions for precision medicine, but these methods have limitations. The results from interaction trees are easy to interpret, but since the covariate-treatment interactions are maximized at each level of the tree, final trees do not connect with any objective function. Thus, it is hard to define an optimal solution for patients. Two-step methods usually need to impose parametric or semi-parametric models to estimate a score function in the first stage and are subject to model misspecification. Zhao et al. (2012) used a weighted classification framework and support vector machine (SVM) to estimate optimal treatment rules but solved the problem in the dual space. Consequently, results are hard to interpret and adopt for clinicians. In order to widely implement an ITR to inform clinical practice, interpretability of the rule is key. To this end, many value function based

procedures have recently been proposed (Xu et al., 2015; Laber and Zhao, 2015; Zhang et al., 2015; Fu et al., 2016). Among them, tree-based methods offer interpretability and enjoy great popularity among clinicians (Su et al., 2009; Lipkovich et al., 2011; Laber and Zhao, 2015). In this paper, we also focus on the value function approach. We develop a recursive partitioning tree algorithm and random forest procedure to optimize the value function in order to obtain ITR.

The tree method was first proposed by Morgan and Sonquist (1963), advanced by the development of classification and regression trees (CART) (Breiman et al., 1984), and is useful in partitioning the predictor space into meaningful subgroups which may elucidate some underlying structure relating a response to predictors. Since an individual tree is known to be highly variable, random forests (Breiman, 2001), an extension of the bagging procedure (Breiman, 1996), was proposed to improve stability. However, both tree and random forest algorithms are typically used for supervised learning where correct outcome labels are provided. Recently Laber and Zhao (2015) proposed a tree algorithm to optimize the value function and search for ITR. Their method is different from original tree algorithms in that they incorporate an objective function (i.e., value function) along with treatment labels through a “purity” measure. Our proposed method may look similar as Laber and Zhao (2015)’s, but differs in the definition of the purity measure. We seek to maximize the value in a given node, conditional on existing treatment assignments in all other previously defined terminal nodes. Unlike traditional regression or classification tree methods where the error function (e.g., sum of squared error, misclassification error) is the widely adopted purity measure and is additive across all final nodes, the value function is a type of overall average of the outcome. Therefore, optimizing the value function within a single node as proposed by Laber and Zhao (2015) may lead to a suboptimal solution which cannot be translated to the overall optimized value defined using all final nodes. When the decision signal level is relatively weak, a single tree decision rule is highly variable. Our random forest algorithm generates a decision rule by averaging over all decision trees in the forest. The decision rule for a future patient is then a soft probability rather than a hard choice. This feature is greatly needed in clinical practice as the strength of the treatment recommendation allows physicians to make a treatment choice using the decision rule along with their judgment and experience. A variable importance measure similar to the one proposed by Breiman (2001) is also developed as a tool for guiding clinical decision making.

In summary, our contribution in this paper is: (1) development of a value function ensemble tree algorithm, ITR forest, to generate an ITR; (2) ITR treatment probability estimation via an ITR forest; (3) implementation of variable importance measure to guide decision making; (4) demonstration of the methods through simulation, randomized controlled trials (RCT), and electronic medical record (EMR) examples. The software package **ITR.Forest** implemented by the statistical computing language **R** (R Core Team, 2017) along with documentation is available at <https://github.com/kdoub5ha/ITR.Forest>.

2 Methods

2.1 Value Function and Individualized Treatment Rules

We are given a random sample of size N from a large population. For each unit i in the sample, where $i = 1, \dots, N$, let t_i be the treatment assignment, y_i be the response, and $\mathbf{x}_i$ be the $p \times 1$ vector of baseline covariates or markers. $(\mathbf{Y}, \mathbf{T}, \mathbf{X})$ is the generic random variable of $\{(y_i, t_i, \mathbf{x}_i)\}$. We let X_j represent the j th covariate where $j = 1, \dots, p$. We denote the distribution of $\mathcal{P}$ by $(\mathbf{Y}, \mathbf{T}, \mathbf{X})$, with E being the expectation with respect to $\mathcal{P}$. For a binary treatment regime, t_i indicates whether the treatment of interest is received, with $t_i = 1$ indicating that the subject received active treatment and $t_i = 0$ indicating that the subject received control, i.e., $t_i \in \{0, 1\}$. Using the potential outcome notation, let $y_i(0)$ denote the outcome under control and $y_i(1)$ the outcome under treatment. We observe t_i and y_i where $y_i = t_i y_i(1) + (1 - t_i) y_i(0)$. Let $r(\mathbf{X})$ denote a vector of binary ITR that is a function of the subjects' baseline covariates $\mathbf{X}$. For any given individualized treatment recommendation r , we let $\mathcal{P}^r$ be the distribution of $(\mathbf{Y}, \mathbf{T}, \mathbf{X})$ given that $\mathbf{T} = r(\mathbf{X})$. Throughout our paper, we use I as an indicator operator where $I(x)$ takes value 1 or 0 if the scalar x is "true" or "false" respectively. When applied to a vector or matrix, I is an element-wise operator.

The research question for individualized treatment recommendation or subgroup identification is only valid when multiple treatment options are available for the same subject. If only one treatment option is allowed or available for certain subjects, the optimal treatment is the only available one. Therefore, without loss of generality, our population space Ω is defined as $\Omega = \{\mathbf{X} \mid \Pr(\mathbf{T} \in (0, 1)) > 0, \forall \mathbf{T} \in \{0, 1\}\}$. Since

$d\mathcal{P} = \Pr(\mathbf{Y} \mid \mathbf{X}, \mathbf{T}) \Pr(\mathbf{T} \mid \mathbf{X}) \Pr(\mathbf{X})$ and $d\mathcal{P}^r = \Pr(\mathbf{Y} \mid \mathbf{X}, \mathbf{T}) I(\mathbf{T} = r(\mathbf{X})) \Pr(\mathbf{X})$, we have,

$\frac{d\mathcal{P}^r}{d\mathcal{P}} = \frac{I(\mathbf{T} = r(\mathbf{X}))}{\Pr(\mathbf{T} \mid \mathbf{X})}$. The expected value of treatment benefit with respect to r is,

$$V(r) = E^r(\mathbf{Y}) = \int \mathbf{Y} d\mathcal{P}^r = \int \mathbf{Y} \frac{d\mathcal{P}^r}{d\mathcal{P}} d\mathcal{P} = E \left[\frac{I(\mathbf{T} = r(\mathbf{X}))}{\Pr(\mathbf{T} \mid \mathbf{X})} \mathbf{Y} \right]. \quad (1)$$

Our goal is to estimate r_o such that, $r_o \in \arg \max V(r)$ (Qian and Murphy, 2011; Zhao et al., 2012; Laber and Zhao, 2015; Zhang et al., 2012a). Using double expectation rule we have,

$$\begin{aligned} V(r) &= E \left[\frac{I(\mathbf{T} = r(\mathbf{X}))}{\Pr(\mathbf{T} \mid \mathbf{X})} \mathbf{Y} \right] = E \left[E \left[\frac{I(\mathbf{T} = r(\mathbf{X}))}{\Pr(\mathbf{T} \mid \mathbf{X})} \mathbf{Y} \mid \mathbf{T}, \mathbf{X} \right] \right] \\ &= E[I(r(\mathbf{X}) = 1) \{E(\mathbf{Y} \mid \mathbf{T} = 1, \mathbf{X}) - E(\mathbf{Y} \mid \mathbf{T} = 0, \mathbf{X})\}] + \\ &\quad E\{E(\mathbf{Y} \mid \mathbf{T} = 0, \mathbf{X})\}. \end{aligned}$$

Therefore our optimizer maximizing $V(r)$ with respect to r is

$$r_o(\mathbf{X}) = I(E(\mathbf{Y} \mid \mathbf{T} = 1, \mathbf{X}) > E(\mathbf{Y} \mid \mathbf{T} = 0, \mathbf{X})). \quad (2)$$

The interpretation of equation (2) is straightforward: simply assign a treatment to patients who can benefit more from it. Equation (2) also connects our method with other personalized treatment suggestion methods such as the contrast function $D(\mathbf{X}) = E(\mathbf{Y} \mid \mathbf{T} = 1, \mathbf{X}) - E(\mathbf{Y} \mid \mathbf{T} = 0, \mathbf{X})$ that was defined and estimated by Cai et al. (2011) and Foster et al.

(2011) to generate ITR. In the context of precision medicine, people often assume that the responses Y are from the following model $Y = \beta_0 + g(X) + T \times D(X) + \epsilon$, where β_0 is the overall mean, both $g(X)$ and $D(X)$ are functions of baseline markers and centered at 0. Based on equation (2), it is easy to see that the optimal solution is, $r_o(X) = I(D(X) > 0)$. Therefore, our method also targets the treatment by marker interactions, similar to the interaction tree approaches (Su et al., 2009; Lipkovich et al., 2011).

With multivariate X , using fully non-parametric methods and incorporating non-linear functions for approximating $r_o(X)$ are extremely valuable, especially when $D(X)$ takes a complex form (Foster et al., 2011). However, these methods are subject to the curse of dimensionality and pose challenges in making inferences about the resulting ITR and its associated value function. On the other hand, if $D(X)$ is estimated by imposing parametric or semi-parametric models on $E(Y|X, T)$, the plug-in estimate of $r_o(X)$ may lead to a much lower population average outcome compared to that of the true $r_o(X)$ (Qian and Murphy, 2011). One may reduce model misspecification by including non-linear bases and selecting important variables via regularized estimation (Qian and Murphy, 2011; Imai and Strauss, 2010). However, it remains challenging to efficiently choose non-linear basis functions to achieve an optimal bias and variance trade-off.

2.2 ITR Tree and Forest

We propose a nonparametric approach that maximizes the value function using a recursive partitioning algorithm and extend this method to generate random forests. We consider a study designed to assess a binary treatment effect on a continuous response with a number of baseline covariates present, either continuous or categorical. Our goal is to search for a treatment assignment with maximized value function. The value function is evaluated by the observed data and is estimated by,

$$\hat{V}(r) = \left(\sum_{i=1}^N \frac{I(t_i = r(x_i))}{\hat{\Pr}(t_i | x_i)} y_i \right) / \left(\sum_{i=1}^N \frac{I(t_i = r(x_i))}{\hat{\Pr}(t_i | x_i)} \right) \quad (3)$$

where $\hat{\Pr}(t_i | x_i)$ is an estimated propensity score and is typically estimated using logistic regression in observational studies (Rosenbaum and Rubin, 1983). This estimator of the value is referred to as the inverse probability weighted estimator (IPWE). We also incorporate the augmented inverse probability weighted estimator (AIPWE) used by Zhang et al. (2012b) in our investigation. The AIPWE protects against misspecification of $\hat{\Pr}(t_i | x_i)$. For RCT, $\hat{\Pr}(t_i | x_i) = 0.5$, assuming a 1 : 1 allocation ratio.

We now introduce how to grow a tree.

- (1) **Initial Split.** For a given space $\Omega = \{x_i, i = 1, \dots, N \mid \Pr(t | x_i) \in (0, 1), \forall t \in \{0, 1\}\} \subset \mathbb{R}^p$ or a single node tree containing all patients, we start with an initial split. This split is induced by a threshold, c , on a covariate, e.g., X_1 . If X_1 is continuous, then the binary decision $r(X)$ is considered as $r(X) = I(X_1 > c)$. If X_1 is nominal with d distinct categories $C = \{c_1, \dots, c_d\}$, then the binary question becomes “Is $X_1 \in A$?” for any subset $A \subset C$. Observations answering “yes” go to the left child node and observations

answering “no” go to the right child node. We assign the new treatment label to the subject based on the child node in which that subject is placed. For example, if the subject goes to left, we consider the treatment decision as treated, otherwise it is control. We define this process by a partition:

$$\{\Omega^1 = \{\Omega \mid X_j \leq c\}, r(\Omega^1) = t\} \quad \text{and} \quad \{\Omega^2 = \{\Omega \mid X_j > c\}, r(\Omega^2) = t'\}, \quad t, t' \in \{0, 1\},$$

where $r(\Omega) = t$ is defined as $r(x_j) = t, \forall x_j \in \Omega$. The value function $\hat{V}(r)$ is then reevaluated using original treatment assignment and the new assignment

$$\hat{V}(j, c, t, t') = \left(\sum_{i \in \Omega^1} \frac{I(t_i = t)}{\hat{\Pr}(t_i \mid x_i)} y_i + \sum_{i \in \Omega^2} \frac{I(t_i = t')}{\hat{\Pr}(t_i \mid x_i)} y_i \right) / \left(\sum_{i \in \Omega^1} \frac{I(t_i = t)}{\hat{\Pr}(t_i \mid x_i)} + \sum_{i \in \Omega^2} \frac{I(t_i = t')}{\hat{\Pr}(t_i \mid x_i)} \right).$$

The best split of the variable X_j is the one that yields the maximum value function among all permissible splits of all markers as well as new treatment assignment, that is, given a covariate X_j and a split point c , we evaluate value function when $r(\Omega_1) = 1$ and $r(\Omega_2) = 0$ versus $r(\Omega_1) = 0$ and $r(\Omega_2) = 1$, i.e.,

$$\hat{V}^* = \max_{j, c, (t, t') \in \{(0, 1), (1, 0)\}} \hat{V}(j, c, t, t').$$

After the initial split the space is partitioned into two subspaces or nodes, Ω^1 and Ω^2 , each with a new treatment label noted as $\hat{r}_o(\Omega^1)$ and $\hat{r}_o(\Omega^2)$, and an estimated maximum value $\hat{V}^*$ for current tree.

- (2) **Second Split.** We now proceed to the next split of both nodes generated from initial step, ω^1 and Ω^2 . For each node, without loss of generality, we denote it as ω^1 . Starting with all of the data in Ω^1 , consider a splitting variable j , split point c , and define a new partition in subspace Ω^1 as, $\Omega^{11} = \Omega^1 \cap \{\Omega \mid X_j \leq c\}$ and $\Omega^{12} = \Omega^1 \cap \{\Omega \mid X_j > c\}$. We seek the splitting variable j , split point c , and treatment assignment $(t, t') \in \{(0, 1), (1, 0)\}$ that solves

$$\max_{j, c, (t, t') \in \{(0, 1), (1, 0)\}} \frac{\sum_{i \in \Omega^{11}} \frac{I(t_i = t)}{\hat{\Pr}(t_i \mid x_i)} y_i + \sum_{i \in \Omega^{12}} \frac{I(t_i = t')}{\hat{\Pr}(t_i \mid x_i)} y_i + \sum_{i \in \Omega^2} \frac{I(t_i = \hat{r}_o(x_i))}{\hat{\Pr}(t_i \mid X_i)} y_i}{\sum_{i \in \Omega^{11}} \frac{I(t_i = t)}{\hat{\Pr}(t_i \mid x_i)} + \sum_{i \in \Omega^{12}} \frac{I(t_i = t')}{\hat{\Pr}(t_i \mid x_i)} + \sum_{i \in \Omega^2} \frac{I(t_i = \hat{r}_o(x_i))}{\hat{\Pr}(t_i \mid X_i)}},$$

and $\hat{V}^{(*)}$ represents the maximized value of above function. This split will only be possible if $\hat{V}^{(*)}$ is greater than $\hat{V}^*$ by a pre-defined level. The same procedure applies to splitting Ω^2 . We adopt a sequential splitting procedure where the order in which nodes Ω^1 and Ω^2 are split is random (Dusseldorp and Van Mechelen, 2014). Finally, we update partition of the space and the treatment assignment $\hat{r}_o$ for each subspace/node, as well as the estimated maximum value $\hat{V}^*$ for current tree.

- (3) Further Splits. At each step of the algorithm, we consider all potential splits of every terminal node of the tree. We then repeat the previous procedure shown in the second split. The order of updating terminal nodes is random. Again, for every node we retain a split if the overall value function is increased by a certain level.

The algorithm stops when a split can no longer be found to increase the value by a pre-defined threshold. In constructing the initial tree, a terminal node is declared when any one of the following conditions is met: (1) the total number of observations in the node is less than some preset minimum node size; (2) the depth of the node is greater than some preset maximum tree depth.

Our partitioning algorithm is a sequential algorithm as noted by Dusseldorp and Van Mechelen (2014). However, such order effects can be counterbalanced by random forest algorithm when a random pre-selection of variables is employed (Strobl et al., 2009). The ITR is developed by maximizing an overall value function, such that in determining c^* we calculate values based on all terminal nodes and their treatment assignments. This differentiates our method from that used by Laber and Zhao (2015). The algorithm stays the same for categorical covariates when subspaces are defined by $\Omega^1 = \{\Omega | \mathbf{X}_1 \in A\}$ and $\Omega^2 = \{\Omega | \mathbf{X}_1 \notin A\}$.

A random forest of ITR trees is then generated via bootstrap sampling, often referred to as “bagging” (Breiman, 2001). Bagging averages many noisy but approximately unbiased models and hence reduces the variance. From a bootstrap sample, an ITR tree is grown with each split determined by maximizing the value function of a randomly selected subset of m_0 predictors. For a set of p predictors we set $m_0 = \max(\lfloor p/3 \rfloor, 1)$ as suggested by Friedman et al. (2001). Repeating this procedure yields a random forest of uncorrelated ITR trees. Our final recommended ITR for each subject is the average of decisions over all trees in the ITR forest. A forest of J trees is denoted by $\mathcal{F} = \{\tau_j : j = 1, 2, \dots, J\}$, where each τ_j is a single tree. For an individual i , a single ITR tree votes for treatment $r_{\tau_j}(\mathbf{x}_i) = 1$ or for control $r_{\tau_j}(\mathbf{x}_i) = 0$. Treatment assignment based on forest $\mathcal{F}$ is calculated as

$p_{\mathcal{F}} = \hat{\Pr}(r_{\mathcal{F}}(\mathbf{x}_i) = 1) = \frac{1}{J} \sum_{j=1}^J r_{\tau_j}(\mathbf{x}_i)$, which is leveraged to estimate the value function for a forest as,

$$\hat{V}_{forest(r)} = \frac{\sum_{i=1}^N \frac{I(t_i = r(\mathbf{x}_i)) [p_{\mathcal{F}} I(t_i = 1) + (1 - p_{\mathcal{F}}) I(t_i = 0)]}{\hat{\Pr}(t_i | \mathbf{x}_i)} y_i}{\sum_{i=1}^N \frac{I(t_i = r(\mathbf{x}_i)) [p_{\mathcal{F}} I(t_i = 1) + (1 - p_{\mathcal{F}}) I(t_i = 0)]}{\hat{\Pr}(t_i | \mathbf{x}_i)}}.$$

It was previously noted by Fu et al. (2016) and Laber and Zhao (2015) that the variance of the value estimated by equation (3) may become unstable when using the raw outcome $\mathbf{Y}$. As a hedge against this instability residuals from a model may be used as the outcome measure. This gives $\mathbf{Y}^* = \mathbf{Y} - \hat{m}(\mathbf{X})$ where $\hat{m}(\mathbf{X})$ are predicted values from some model. Fu et al. (2016) and Laber and Zhao (2015) recommended linear models and random forests to estimate $\hat{m}(\mathbf{X})$, respectively. For the remainder of this paper we use random forests to estimate $\hat{m}(\mathbf{X})$.

unless otherwise noted. We utilize the R package **randomForest** with the defaults to estimate $\hat{m}(X)$ when random forests are used to stabilize the variance (Liaw and Wiener, 2002). Y^* is used to generate an ITR in training data and original outcome Y is used to calculate maximized value by validation data.

2.3 Pruning

A large tree likely suffers from overfitting as the tree growing algorithm is greedy. Breiman et al. (1984) suggested a pruning procedure that creates a set of optimally pruned subtrees of Γ_0 , and selects the best subtree based on a tree complexity penalty. The performance of an ITR tree is evaluated using

$$V_\lambda(\Gamma) = V(\Gamma) - \lambda \cdot |\Gamma - \tilde{\Gamma}|, \quad (4)$$

where $V(\Gamma)$ evaluates the total value in tree Γ . $V(\Gamma)$ is obtained by evaluating equation (3) for the decision rule generated by tree Γ . $\tilde{\Gamma}$ is the set of terminal nodes of tree Γ , $|\Gamma - \tilde{\Gamma}|$ is the number of internal nodes in Γ , and $\lambda > 0$ is a penalty based on tree size. A tree with larger $V_\lambda(\Gamma)$ is desirable and a cross validation procedure is used to select tuning parameter λ .

Given a training data set and a penalty parameter λ , a large initial tree Γ_0 and calculate $V_\lambda^{train}(\Gamma)$. Next, we want to trim the “weakest” branch of Γ_0 . To find the “weakest” branch of Γ_0 , consider each non-terminal node h of Γ_0 and the branch Γ_h which has h as its root node. The “weakest” branch of Γ_0 satisfies the following conditions: (1) node h has the greatest number of parent nodes (i.e. is far away from the initial node of Γ_0) and (2) $\Gamma_0 - \Gamma_h$ has the highest value among all branches Γ_h . $V_\lambda^{train}(\Gamma_0 - \Gamma_h)$ is the overall value of the initial tree after branch Γ_h is trimmed. The “weakest” branch is both the farthest away from the root node of Γ_0 and also contributes the smallest additional value to the tree. Hence, if two branches, say Γ_{h_1} and Γ_{h_2} , have root nodes equidistant from the root node of Γ_0 , we would prune Γ_{h_2} which satisfies $\arg\max_i \{V_\lambda^{train}(\Gamma_0 - \Gamma_{h_i})\}$. This results in subtree $\Gamma_1 < \Gamma_0$ where $<$ means “subtree of”. Repeating this procedure until the tree is pruned back to the root node results in a series of optimally pruned subtrees $\Gamma_M < \Gamma_{M-1} < \dots < \Gamma_1 < \Gamma_0$ with Γ_M being the null tree having no splits. The optimal subtree will be the subtree which maximizes equation (4), denoted as Γ_m . The test sample is then run down Γ_m to obtain the validated value $V_\lambda^{test}(\Gamma_m)$. We repeat this procedure for every partitioned data pair (e.g. 5-fold) and calculate the mean validated value for the given λ . The final λ is chosen so that mean validated value is maximized over a sequence of values of λ . A similar pruning procedure was used by Su et al. (2009).

2.4 Variable Importance Measure

When there are new markers introduced to assist in treatment selection, it is important to evaluate their value in improving population average outcomes, in our case ITR. We designed a variable importance measure in order to evaluate which markers are important for an ITR. Our variable importance algorithm is based on random forests (Breiman, 2001) and

was similarly used in the interaction tree method (Su et al., 2009). First, a random forest of ITR trees is constructed. For a bootstrap sample L_b , each tree Γ_b represents an ITR. Next, the sample not in L_b , $L_b^\perp$, is sent down the tree and the value $V(\Gamma_b)$ computed as defined in equation (3). After permuting the j th marker in $L_b^\perp$, the permuted sample is evaluated by Γ_b and $V_j(\Gamma_b)$ is recomputed. Finally, the relative difference between $V(\Gamma_b)$ and $V_j(\Gamma_b)$ is recorded. This procedure is repeated for B bootstrap samples. We define the importance measure W_j to be the average of relative differences over B bootstrap samples, i.e., $W_j = \frac{1}{B} \sum_b \frac{V(\Gamma_b) - V_j(\Gamma_b)}{V(\Gamma_b)}$. Further details on the algorithm can be found in Supplementary Material Algorithm 2 and Breiman (2001).

3 Simulation

This section contains simulated experiments designed to investigate the performance of the ITR forest and tree procedures. Data is generated from the models outlined in Table 1. Each data set consists of a continuous response Y , a binary treatment T , and four independent covariates X_1, X_2, X_3, X_4 from a unit uniform distribution. Unless otherwise stated, simulation replicates also include 10 excess noise predictors from a unit uniform distribution.

Simulation A is an RCT design while simulations B - D are observational study designs. Simulation models A and B have a single tree structure. Models A.1 and B.1 produce a tree with one split and two terminal nodes: an initial split is at $X_1 = 0.5$ with the treatment sent to those with $X_1 < 0.5$. Models A.2 and B.2 produce a tree with two splits and three terminal nodes. The initial split is at $X_1 = 0.3$ and the right child node is further split at $X_3 = 0.1$. Treatment is assigned to patients with $X_1 > 0.3$ and $X_3 > 0.1$. Models C.1 and C.2 have well defined subgroups which benefit from treatment, but the ITR cannot be defined by a tree structure. In model C.1, all patients benefit from treatment with increasing values of X_1 accompanied by an increase in value defined by equation (3). In model C.2, patients with $X_1 + X_3 > 1$ benefit from treatment and those with greater values of X_2 or X_4 receive increased benefit. In both C schemes, treatment effects are smooth functions of covariates. Scheme D is a null model with no defined subgroup and should produce a tree with no splits. We use accuracies of correctly predicting treatment assignment given an ITR to evaluate the performance of ITR trees and forests. Maximized value is also used to evaluate the procedure, particularly in comparison with other methods (Matsouaka et al., 2014). For simulations A and B, signal ratio is defined as $\theta = \beta_1/\beta_2$. Signal ratio compares the benefit of receiving treatment for those in the subgroup (β_1) with the benefit of receiving control for those not in the subgroup (β_2). We fix $\beta_2 = 1$, and vary β_1 so that $\theta \in \{1, 3, 5\}$. For simulation C, the linear interaction effect β is varied from 1 to 5. For all simulations, models were trained on $n = 300, 500$ and 800 observations and validated on $n = 1000$ observations. All simulation analyses used 100 replicates.

3.1 Pruning

Accuracy of the pruning procedure is reported in Table 2. It contains the percent of correct sized trees after pruning and a summary of the selected penalty parameter λ values. The

penalty λ was selected using the method presented in section 2.3 with a 5-fold cross validated estimate of λ . Values of λ vary depending on sample size, effect size, and outcome variable as previously noted by (Laber and Zhao, 2015). Performance is notably better when $\theta = 1$ versus $\theta = 3$ for simulations A.2 and B.2. This is likely due to the treatment effect differential between those who benefit from treatment versus those who benefit from control. While stabilization of the variance may address some of the issues associated with this imbalance, a single tree structure may not be powerful enough to overcome the effect difference. In general, accuracies increase as sample size increases. For training samples of $n = 800$ the correct sized tree was grown at least 78% of the time, with the exception of simulation scheme B.2. When $\theta = 1$ a training sample of $n = 800$ grew the correct sized tree in 67% of replicates. Scheme D is used to estimate the type I error rate. The null tree was selected in at least 99% of replicates for scheme D.

A larger sample size example with $n = 2500$ observations from scheme B.2 with $\theta = 3$ is reported (Figure 1). It is used so that a large initial tree structure is generated. The optimal $\lambda = 0.04$ was selected using 5-fold cross validation. The resulting optimal tree structure selects both correct splitting variables and correct cut points and contains no superfluous splits. Further details related to this example can be found in the documentation of our software package (<https://github.com/kdoub5ha/ITR.Forest>).

3.2 Accuracy

Accuracy of the ITR forest and tree methods is assessed for simulations A-C. The choice of splitting variable and cut point is summarized in the context of our pruning procedure by generating a large tree, pruning the tree (per section 2.3), and using a validation sample of $n = 1000$ observations to assess whether the correct treatment is assigned to a patient. Figure 2 presents the average and the standard deviation of the ITR tree procedure accuracies over 100 simulation replicates. Signal ratio, $\theta = 1, 3, 5$ or linear interaction effect, $\beta = 1, 3, 5$ and sample sizes are $n = 300, 500, 800$. Table 3 shows results from constructing a single ITR forest. We generated a single data set of $n = 300$ observations with $\theta, \beta = 3$ for each simulation scheme, and each data set was fit with an ITR forest of 500 trees and an ITR tree model with 5-fold cross validated pruning. Treatment predictions from each model were obtained for a validation sample of $n = 1000$ observations. The top panel in Table 3 shows the patient level forest summary, the probability of being assigned to active treatment using all trees in the ITR forest, for a randomly selected subset of 10 validation individuals. IPWE is shown for all schemes and AIPWE results are shown for schemes B and C. For ITR forest and tree methods random forests were used for variance stabilization. A linear model was initially used for variance stabilization in scheme C, but ITR tree methods both produced null models so random forests were used instead.

A probability greater than 0.5 indicates a majority of trees voted for active treatment, and active treatment would be recommended for this individual. Otherwise, control would be recommended. If the probability is close to 0.5, a treatment recommendation may be withheld. We note that for a 500 tree ITR forest the standard deviation for treatment assignment of 0.5 is $0.5 / \sqrt{500} = 0.022$. Hence, when a treatment assignment probability between 0.478 and 0.522, the decision may be made in combination with a clinician's

experience. Subject 5 from scheme A.1, for instance, has a predicted probability of benefiting from active treatment of 0.49. Although this gives the optimal treatment recommendation for this patient of being on control, the expertise of the patient's physician should come into play as well. This subject has covariate values $X_1 = 0.04$ and $X_3 = 0.94$, so that they satisfy the subgroup inclusion criteria strongly for X_3 but not for X_1 . In a few instances a subject received an incorrect treatment recommendation. In scheme B.2, observations with $X_1 > 0.3$ and $X_3 > 0.1$ should be assigned to active treatment, otherwise control. Subject 1 received an incorrect recommendation to be on active treatment using both IPWE and AIPWE methods. This subject had $X_1 = 0.72$ and $X_3 = 0.02$ and so strongly meets the subgroup criteria for X_1 , but not for X_3 . In scheme C.1, all 10 subjects were recommended to be on active treatment as was expected. In scheme C.2, subject 10 was misclassified. This subject had $X_1 = 0.94$ and $X_3 = 0.16$ so that the sum of these two covariates is close to 1, leading to the misclassification. The bottom panel of Table 3 shows the percent of all 1000 validation observations from the ITR forest and ITR tree models which received correct treatment assignments. A single ITR tree structure performed well compared to the ITR forest. The ITR forest, however, has the additional benefit of the soft treatment assignment probability and this is made clear in Table 4. The accuracy for RCT simulations (A) was better than EMR simulations (B) due to the need to estimate $\hat{\Pr}(t_i | x_i)$. Additionally, we investigated use of training samples of $n = 100$, but found the smaller sample size prohibitive to discovery of subgroups. Thus, only sample sizes of $n = 300$ or greater are presented in this paper.

3.3 Variable Importance

Variable importance measures are reported for simulated data with sample sizes of $n = 300$, 500, and 800. All variable importance measures were estimated using a 500 tree ITR forest, and the reported variable importance measures were scaled so that $\sum_j W_j = 1$ for easier interpretation. We note that when excess noise predictors were included in estimating variable importance, the excess predictors received at most 4% of the importance measure and typically received less than 0.1%. Hence, we exclude noise variables from variable importance analysis in this paper. A summary of variable importance measures for simulated data is presented in Figure 3. Note that for simulations A.1, B.1, and C.1 the subgroup is defined by X_1 and for simulations A.2, B.2, C.2 the subgroup is defined by X_1 and X_3 . IPWE was used for scheme A while AIPWE was used for schemes B and C. The ITR forest procedure showed excellent ability to select the correct subgroup defining variable(s) as most important. Since the cut point in schemes A.2 and B.2 of $X_1 = 0.3$ is farther from the uniform (0,1) boundary than the cut point for X_3 of 0.1, the variable importance measure for X_1 was greater than X_3 . This is because there are fewer observations sampled on average with $X_3 < 0.1$ than with $X_1 < 0.3$. This point is illustrated in the supplementary material (Figure S1) by considering sampling schemes with both cut points set at 0.3. If both cut points are 0.3 the variable importance measures for X_1 and X_3 were similar. Scheme C.1 showed X_1 as most important predictor once $n = 500$ and the effect size $\beta = 2$. Scheme C.2 showed most of the importance measures going to X_1 and X_3 , even for smaller training samples.

3.4 Method Comparison

The maximized value of the ITR forest and tree methods was compared to a few optimal treatment recommendation methods. Methods considered were interaction trees (Su et al., 2009) (IT tree), outcome weighted learning using SVM (Zhao et al., 2012), and minimal node impurity measure (MIDAs) (Laber and Zhao, 2015). ITR trees were pruned using 5-fold cross validation. Optimal tree size selection for MIDAs was similar, but used a 10-fold cross validation per the suggestion in Laber and Zhao (2015). Random forests were used to stabilize variance in MIDAs. Bayesian Information Criterion (BIC) was used in IT tree methods along with the amalgamation algorithm. Outcome weighed learning via SVM used a linear kernel and 5-fold cross validation to find an ITR in R package **DTRlearn** (Liu et al., 2015). We also include random guessing as a reference for arbitrary treatment assignment. Each model was trained on $n = 300$ observations with $\theta = 3$ or $\beta = 3$ and validated on $n = 1000$ observations. IPWE was used for simulation A.2 and AIPWE used for the other simulations. Results are found in Table 4. Simulations A.1 and B.1 were excluded from this section of the paper since the ITR tree and MIDAs methods are identical estimators of rewards/value defined by equation (3) when the tree structure consists of only a single split. Different number of excess prognostic markers (noise) were included, generated from a uniform distribution. Results for the null model D were consistent for all methods and random guessing resulted in the smallest value. ITR forest gave the best estimated maximized value for all simulations considered, demonstrating the utility of the ITR forest method and its robustness.

4 Data Applications

The ITR forest and tree methods were applied to RCT data from the DURAbility of Basal versus Lispro mix 75/25 insulin Efficacy (DURABLE) study (Buse et al., 2009) and an EMR data set from Clinical Practice Research Datalink (CPRD) (<https://www.cprd.com/intro.asp>). DURABLE investigated two treatments, a mix of 75% insulin lispro protamine suspension and 25% lispro (LM75/25) vs Glargine, in patients with type 2 diabetes with the objective of achieving glycemic control. CPRD is an observational and intervention based research service that operates as part of the UK Department of Health. EMR data was obtained by Eli Lilly and Company in order to compare two injectable treatments, basal insulin and glucagon-like peptide-1 receptor agonists (GLP-1), in patients with type 2 diabetes in an effort to control glycemic change (Lawson et al., 1998). The outcome used for these models is percent change in HbA1c from baseline. Optimal tree size was determined using the pruning procedure outlined in section 2.3 using a 5-fold cross validation estimator of λ . A random sample of 10 subjects were set aside in each study to assess the model performance and the remaining individuals used to train the model. A 500 tree ITR forest was constructed using all available variables. We compared the ITR generated by both tree and random forest using the 10 held-out individuals. Random forests were used to stabilize the variance for both data sets.

4.1 DURABLE study

For the DURABLE study, the variables available at baseline were fasting plasma glucose, insulin, 7-points glucose readings taken throughout a day, weight, height, BMI (body mass

index), diastolic blood pressure, systolic blood pressure, heart rate, and duration of diabetes in $n = 1385$ patients. Of the 1385 patients, 688 received LM75/25 and 697 received Glargine as control. After setting aside 10 subjects for validation, the remaining 1375 observations were used to construct an ITR forest and an ITR tree. The variable importance measure returned weight as the most important predictor of treatment assignment using the IPWE and AIPWE (Figure 5). A tree with 2 terminal nodes was determined to be optimal for both IPWE and AIPWE (Figure 4). There was correspondence between patient weight receiving a plurality of the importance measure and being the covariate selected for the initial split. Table 5 shows treatment assignment using an ITR forest. Patients 1, 6, and 9 were interesting cases, receiving opposing decisions from the ITR forest and ITR tree models. These three patients had weights of 115 kg, 106 kg, and 129 kg and so were recommended to take Glargine by the ITR tree. These results were consistent with the mechanism of reactions. Glargine is a basal insulin which lowers the overall blood glucose, and LM 75/25 is a pre-mixed insulin which covers both postprandial and basal glucose. Heavier patients need more basal insulin and LM 75/25 may not deliver enough without causing hypoglycemia. Therefore, Glargine works better for these types of patients.

4.2 Electronic Medical Records

The EMR data set contained 5177 patient observations, 837 treated with GLP-1 and 4340 with basal insulin. Variables available at baseline were the indexed laboratory result, age, diabetes duration, Charlson score, ethnicity, race, categorical body mass index, and insurance status. The lab result is a measurement of a patient's average blood glucose level during the past 3 months and Charlson score is a composite score of patient comorbidities. The pruning procedure selected a tree with two terminal nodes for the IPWE and AIPWE (Figure 4). Index lab result was selected as the splitting variable with patients having a lab result lower than 12 being assigned to GLP-1 and basal insulin otherwise. Treatment assignments are shown in Table 5. Variable importance measures returned the Charlson score as the most important predictor of treatment assignment for both the IPWE and AIPWE (Figure 5). All 10 of the patients held out of the training sample were recommended strongly to be on treatment. All 10 patients had a lab score less than 12 leading to the consistent recommendation of treatment among these patients. These results are helpful for us to better decide on how to initiate first line injectable anti-diabetic treatments when patients have failed on multiple oral medications. Our results indicate that for patients with higher average blood glucose, it is better to start with insulin treatment. For patients with lower average blood glucose, GLP-1 treatment is recommended. The average blood glucose also indicates the disease progression. Patients with lower blood glucose often have better beta cell function where GLP-1 treatments work better.

The computational cost of each method is shown in Table 6 using this EMR data. Each tree using the value function as the reward is grown in less than 25 seconds. The IT tree runs in 1.6 seconds and ITR forest in 38 seconds. OWL is the slowest among all.

5 Discussion

Individualized treatment rules are increasingly being used by clinical and intervention scientists to account for patient response heterogeneity (Ludwig and Weinstein, 2005; Hayes et al., 2007). These treatment rules belong to the new era of precision medicine (Piquette-Miller and Grant, 2007; Hamburg and Collins, 2010). Regression based methods model the response as a function of patient characteristics and treatment, selecting treatments that maximize the predicted mean outcome (Qian and Murphy, 2011). However, because such methods indirectly infer the optimal treatment rule through a regression model, they are subject to model misspecification. Direct methods are an alternative to regression based indirect methods that depend on modeling the conditional mean first. Zhao et al. (2012) demonstrated that optimal treatment rules can be estimated within a weighted classification framework and solved the problem using support vector machines (SVM) via the hinge loss (Cortes and Vapnik, 1995). However, without an interpretable and transparent representation of the model behind these approaches, clinical investigators may be hesitant to use the estimated treatment rule to inform clinical practice or future research. In this paper, we attempted to directly maximize clinical response to estimate optimal individualized treatment rule using a recursive partitioning tree algorithm. The recursive partitioning tree method is a non-parametric search procedure, easy to interpret, and handles high dimensional and large scale modern data sets (e.g., genomics and EMR) seamlessly. We used a random forest ensemble predictor to mitigate the inherent instability of a single tree, which may be a weak standalone predictor.

When there are new biomarkers introduced to assist in treatment selection, it is important to evaluate their value in improving population average outcomes. A variable highly differentially associated with $y_X(1)$ and $y_X(0)$ may not necessarily be important for improving ITRs. This is somewhat similar to the phenomenon observed in the risk prediction literature: a variable highly significant in regression modeling may not result in large improvement in prediction. We then developed variable importance measures based on ITR forests, providing a measure of the overall influence of a covariate in predicting treatment outcome. When the dimension of new markers is large, it would be crucial to employ the cross-validation to correct for the overfitting bias as suggested by Zhao et al. (2012). Procedures for efficiently selecting the informative markers warrant further research.

The decision rule from an ITR tree with only a few splits can make discussion of treatment between patients and health care providers transparent. Through simulation for both RCT and observational designs, we demonstrate the accuracy and stability of our algorithms. The ITR forest outperforms comparable methods in maximizing value function in all the non-trivial scenarios considered in this paper, demonstrating the robustness of the ITR forest. The application to RCT and observational studies further validate the utility of this method. In general, a composition of the ITR tree, forest and variable importance measure will give clinicians useful tools in considering treatment regimens. We recommend that samples be adequately large, so that identification of subgroups, assuming they exist, is feasible.

There are several extensions of our method that can be pursued. Incorporating multiple and continuous treatments into the ITR tree and forest method can be achieved by modifying our

splitting criteria. Sample sizes for most clinical trials are powered for the primary objectives of those studies, and often not for precision medicine or subgroup identifications. Therefore, synthesizing evidence from multiple studies could potentially develop a more robust ITR. Finally, it is known that the greedy search approach induces a bias in variable selection towards variables with more distinct values (Doyle, 1973). Our tree and forest methods can be further calibrated by incorporation of previous efforts to correct this bias (Loh, 2002).

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

ACKNOWLEDGMENTS

The authors gratefully acknowledge two anonymous reviewers for their constructive and helpful comments. KD is supported by University of Arizona University Fellowship. JZ is supported by NIH grant K01DK106116.

References

- Arakaki R, Blevins T, Wise J, Liljenquist D, Jiang H, Jacobson J, Martin S, and Jackson J (2014), "Comparison of insulin lispro protamine suspension versus insulin glargine once daily added to oral antihyperglycaemic medications and exenatide in type 2 diabetes: A prospective randomized open-label trial," *Diabetes, Obesity and Metabolism*, 16, 510–518.
- Bergental RM, Bailey CJ, and Kendall DM (2010), "Type 2 diabetes: Assessing the relative risks and benefits of glucose-lowering medications," *The American Journal of Medicine*, 123, 374–e9.
- Breiman L (1996), "Bagging predictors," *Machine Learning*, 24, 123–140.
- Breiman L, Friedman J, Olshen R, and Stone C (1984), "Classification and Regression Trees CRC Press," Boca Raton, Florida.
- Breiman, v. L. (2001), "Random forests," *Machine Learning*, 45, 5–32.
- Buse JB, Wolffenbuttel BH, Herman WH, Shemonsky NK, Jiang HH, Fahrback JL, Scism-Bacon JL, and Martin SA (2009), "DURABILITY of basal versus lispro mix 75/25 insulin efficacy (DURABLE) trial 24-week results," *Diabetes Care*, 32, 1007–1013. [PubMed: 19336625]
- Cai T, Tian L, Wong PH, and Wei L (2011), "Analysis of randomized comparative clinical trial data for personalized treatment selections," *Biostatistics*, 12, 270–282. [PubMed: 20876663]
- Cortes C and Vapnik V (1995), "Support-vector networks," *Machine Learning*, 20, 273–297.
- Cummings J, Emre M, Aarsland D, Tekin S, Dronamraju N, and Lane R (2010), "Effects of rivastigmine in Alzheimer's disease patients with and without hallucinations," *Journal of Alzheimer's Disease*, 20, 301–311.
- Diamant M, Van Gaal L, Guerci B, Stranks S, Han J, Malloy J, Boardman MK, and Trautmann ME (2014), "Exenatide once weekly versus insulin glargine for type 2 diabetes (DURATION-3): 3-year results of an open-label randomised trial," *The Lancet Diabetes & Endocrinology*, 2, 464–473. [PubMed: 24731672]
- Doyle P (1973), "The use of automatic interaction detector and similar search procedures," *Operational Research Quarterly*, 465–467
- Dusseldorp E and Van Mechelen I (2014), "Qualitative interaction trees: a tool to identify qualitative treatment-subgroup interactions," *Statistics in medicine*, 33, 219–237. [PubMed: 23922224]
- Faries DE, Chen Y, Lipkovich I, Zagar A, Liu X, and Obenchain RL (2013), "Local control for identifying subgroups of interest in observational research: Persistence of treatment for major depressive disorder," *International Journal of Methods in Psychiatric Research*, 22, 185–194. [PubMed: 23956114]
- Forst T, Guthrie R, Goldenberg R, Yee J, Vijapurkar U, Meininger G, and Stein P (2014), "Efficacy and safety of canagliflozin over 52 weeks in patients with type 2 diabetes on background metformin and pioglitazone," *Diabetes, Obesity and Metabolism*, 16, 467–477.

- Foster JC, Taylor JM, and Ruberg SJ (2011), “Subgroup identification from randomized clinical trial data,” *Statistics in Medicine*, 30, 2867–2880. [PubMed: 21815180]
- Friedman J, Hastie T, and Tibshirani R (2001), *The elements of statistical learning*, vol. 1, Springer series in statistics Springer, Berlin.
- Fu H, Zhou J, and Faries DE (2016), “Estimating optimal treatment regimes via subgroup identification in randomized control trials and observational studies,” *Statistics in Medicine*.
- Hamburg MA and Collins FS (2010), “The path to personalized medicine,” *New England Journal of Medicine*, 363, 301–304. [PubMed: 20551152]
- Hayes DF, Thor AD, Dressler LG, Weaver D, Edgerton S, Cowan D, Broadwater G, Goldstein LJ, Martino S, Ingle JN, et al. (2007), “HER2 and response to paclitaxel in node-positive breast cancer,” *New England Journal of Medicine*, 357, 1496–1506. [PubMed: 17928597]
- Imai K and Strauss A (2010), “Estimation of heterogeneous treatment effects from randomized experiments, with application to the optimal planning of the get-out-the-vote campaign,” *Political Analysis*, mpq035.
- Laber E and Zhao Y (2015), “Tree-based methods for individualized treatment regimes,” *Biometrika*, 102, 501–514. [PubMed: 26893526]
- Lawson D, Sherman V, and Hollowell J (1998), “The General Practice Research Database. Scientific and Ethical Advisory Group.” *QJM: An International Journal of Medicine*, 91, 445–452. [PubMed: 9709463]
- Liaw A and Wiener M (2002), “Classification and Regression by randomForest,” *R News*, 2, 18–22.
- Lipkovich I, Dmitrienko A, Denne J, and Enas G (2011), “Subgroup identification based on differential effect search—A recursive partitioning method for establishing response to treatment in patient subpopulations,” *Statistics in Medicine*, 30, 2601–2621. [PubMed: 21786278]
- Liu Y, Wang Y, and Zeng D (2015), *DTRlearn: Learning Algorithms for Dynamic Treatment Regimes*, r package version 1.2.
- Loh W-Y (2002), “Regression tress with unbiased variable selection and interaction detection,” *Statistica Sinica*, 361–386.
- Ludwig JA and Weinstein JN (2005), “Biomarkers in cancer staging, prognosis and treatment selection,” *Nature Reviews Cancer*, 5, 845–856. [PubMed: 16239904]
- Matsouaka RA, Li J, and Cai T (2014), “Evaluating marker-guided treatment selection strategies,” *Biometrics*, 70, 489–499. [PubMed: 24779731]
- Morgan JN and Sonquist JA (1963), “Problems in the analysis of survey data, and a proposal,” *Journal of the American Statistical Association*, 58, 415–434.
- Nyenwe EA, Jerkins TW, Umpierrez GE, and Kitabchi AE (2011), “Management of type 2 diabetes: Evolving strategies for the treatment of patients with type 2 diabetes,” *Metabolism*, 60, 1–23. [PubMed: 21134520]
- Piquette-Miller M and Grant D (2007), “The art and science of personalized medicine,” *Clinical Pharmacology & Therapeutics*, 81, 311–315. [PubMed: 17339856]
- Qian M and Murphy SA (2011), “Performance guarantees for individualized treatment rules,” *Annals of Statistics*, 39, 1180. [PubMed: 21666835]
- R Core Team (2017), *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria.
- Rosenbaum PR and Rubin DB (1983), “The central role of the propensity score in observational studies for causal effects,” *Biometrika*, 70, 41–55.
- Strobl C, Malley J, and Tutz G (2009), “An introduction to recursive partitioning: rationale, application, and characteristics of classification and regression trees, bagging, and random forests,” *Psychological methods*, 14, 323. [PubMed: 19968396]
- Su X, Kang J, Fan J, Levine RA, and Yan X (2012), “Facilitating score and causal inference trees for large observational studies,” *Journal of Machine Learning Research*, 13, 2955–2994.
- Su X, Tsai C-L, Wang H, Nickerson DM, and Li B (2009), “Subgroup analysis via recursive partitioning,” *Journal of Machine Learning Research*, 10, 141–158.
- Xu Y, Yu M, Zhao Y-Q, Li Q, Wang S, and Shao J (2015), “Regularized outcome weighted subgroup identification for differential treatment effects,” *Biometrics*, 71, 645–653. [PubMed: 25962845]

- Zhang B, Tsiatis AA, Davidian M, Zhang M, and Laber E (2012a), “Estimating optimal treatment regimes from a classification perspective,” *Stat*, 1, 103–114. [PubMed: 23645940]
- Zhang B, Tsiatis AA, Laber EB, and Davidian M (2012b), “A robust method for estimating optimal treatment regimes,” *Biometrics*, 68, 4040–4048.
- Zhang Y, Laber EB, Tsiatis A, and Davidian M (2015), “Using decision lists to construct interpretable and parsimonious treatment regimes,” *Biometrics*, 74, 895–904.
- Zhao L, Tian L, Cai T, Claggett B, and Wei L-J (2013), “Effectively selecting a target population for a future comparative study,” *Journal of the American Statistical Association*, 408, 527–539.
- Zhao Y, Zeng D, Rush AJ, and Kosorok MR (2012), “Estimating individualized treatment rules using outcome weighted learning,” *Journal of the American Statistical Association*, 107, 1106–1118. [PubMed: 23630406]

Figure 1:

Example of the pruning procedure in section 2.3 by a simulated example from scheme B.2 with $\theta = 3$, 10 excess noise variables, and a training sample of $n = 2500$ observations. **(Left)** The full tree shown with each split represented by a 3 component vector (splitting variable, splitting value, child node to which active treatment is sent). **(Right)** Plot of 5-fold cross validated penalized value function, $V_\lambda(\Gamma)$, (y-axis) with error bars plotted against the penalty parameter λ (x-axis). The optimally pruned tree structure is embedded in the right panel.

Figure 2:

Percent of validation sample observations assigned to the correct treatment using the ITR tree method. Models were trained on $n = 300$ (circle), 500 (square), and 800 (diamond) observations, pruned using a 5-fold cross validation estimator of penalty parameter λ , and validated using a sample of $n = 1000$ observations. Point estimates and 95% equal tail intervals over 100 replicates are displayed. A.1 and A.2 use IPWE and AIPWE used otherwise. Signal ratios (θ) and linear interaction effect (β) were varied from 1 to 5.

ITR: individualized treatment rule; **IPWE:** inverse probability weighted estimator; **AIPWE:** augmented inverse probability weighted estimator.

Figure 3:

Variable importance measures for simulation studies A, B, and C with $n = 300, 500, 800$.

Signal ratios of $\theta = 1, 3, 5$ (A and B schemes) and $\beta = 1, 3, 5$ (C schemes) using ITR forest method with the IPWE for A schemes and AIPWE for B and C schemes. Gray scale indicates $\theta, \beta = 1$ (black) to $\theta, \beta = 5$ (light gray). The subgroup defining variable(s) for models A.1, B.1, and C.1 is X_1 and for models A.2, B.2, and C.2 are X_1 and X_3 .

ITR: individualized treatment rule; **IPWE:** inverse probability weighted estimator; **AIPWE:** augmented inverse probability weighted estimator

Figure 4:

Estimated ITR for DURABLE RCT (left) and EMR (right) using a single ITR tree structure and optimal pruning as described in section 2.3. RCT model was trained using $n = 1375$ observations and EMR model was trained on $n = 5167$ observations. The IPWE and AIPWE for both data sets returned the same optimal tree. Each internal node is defined by the triple (splitting variable, splitting value, child node to which active treatment, LM75/25 for DURABLE and GLP-1 for EMR, is sent).

ITR: individualized treatment rule; **RCT**: randomized controlled trial; **EMR**: electronic medical records; **IPWE**: inverse probability weighted estimator; **AIPWE**: augmented inverse probability weighted estimator; **DURABLE**: Assessing the DURAbility of basal vs lispro mix 75/25 insulin efficacy trial; **LM75/25**: mix of 75% insulin lispro protamine suspension and 25% lispro (LM75/25); **GLP-1**: glucagon-like peptide-1 receptor agonists.

Figure 5:

Variable importance measures for DURABLE RCT and EMR diabetes data generated using a 500 tree ITR forest with IPWE (black) and AIPWE (gray).

ITR: individualized treatment rule; **RCT:** randomized controlled trial; **EMR:** electronic medical records; **IPWE:** inverse probability weighted estimator; **AIPWE:** augmented inverse probability weighted estimator; **DURABLE:** Assessing the DURability of basal vs lispro mix 75/25 insulin efficacy trial.

Table 1:

Simulation schemes for assessing the performance of ITR tree and forest methods. The A schemes are RCT design while B, C, and D schemes are observational study models. Models A.1 and B.1 have subgroups defined by a single variable and single cut point which should produce a tree with a single split. Models A.2 and B.2 have subgroups defined by two variables and two cut points which should produce a tree with two splits. Models C.1 and C.2 are linear interaction models with one and two interacting variables, respectively. Models C.1 and C.2 cannot be expressed by a single tree. Model D has no subgroups and so should produce a null tree. Covariates X_1, X_2, X_3, X_4 are independent and distributed as $\text{Uniform}(0, 1)$. Errors $\epsilon \sim \mathcal{N}(0, 1)$. Ω represents all samples and $\emptyset$ represents the null set.

Scheme	Model	$\Pr(T X)$	Subgroup A
A.1	$Y = 2 + 2I(X_2 < 0.5) + \beta_1 I(x \in A)T + \beta_2 [1 - I(x \in A)](1 - T) + \epsilon$	0.5	$X_1 < 0.5$
A.2	$Y = 2 + 2I(X_2 < 0.5) + \beta_1 I(x \in A)T + \beta_2 [1 - I(x \in A)](1 - T) + \epsilon$	0.5	$X_1 > 0.3$ and $X_3 > 0.1$
B.1	$Y = 1 + 2X_2 + 4X_4 + \beta_1 I(x \in A)T + \beta_2 [1 - I(x \in A)](1 - T) + \epsilon$	$\text{logit}^{-1}(-4 + 3X_2 + 5X_4)$	$X_1 < 0.5$
B.2	$Y = 1 + 2X_2 + 4X_4 + \beta_1 I(x \in A)T + \beta_2 [1 - I(x \in A)](1 - T) + \epsilon$	$\text{logit}^{-1}(-4 + 3X_2 + 5X_4)$	$X_1 > 0.3$ and $X_3 > 0.1$
C.1	$Y = 6 + 2X_1 + \beta X_1^2 T + \epsilon$	$\text{logit}^{-1}(-4 + 3X_2 + 5X_4)$	Ω
C.2	$Y = 6 + 2T + 2X_1 + 2X_3 + \beta I(X \in A) \exp\{X_2 + X_4\}T + \epsilon$	$\text{logit}^{-1}(-4 + 3X_2 + 5X_4)$	$X_1 + X_3 > 1$
D	$Y = 1 + 2T + 2X_1 + 2X_3 + \epsilon$	$\text{logit}^{-1}(-6 + 3X_1 + 3X_2 + 3X_3 + 3X_4)$	$\emptyset$

ITR: individualized treatment rule; **RCT**: randomized controlled trial.

Table 2:

Proportion (95% interval) of correctly sized trees selected by the pruning procedure over 100 replicates. Penalty parameter λ and its minimum and maximum (in parenthesis) obtained using 5-fold cross validation for the 100 replicates are shown. Schemes A and B have $\theta = 1, 3$ and all schemes presented were trained with samples of $n = 300, 500$, and 800 . All replicates include 10 excess noise variables.

Scheme	n	$\theta = 3$		$\theta = 1$	
		Correct Tree Proportion (95% CI)	λ Mean (Min, Max)	Correct Tree Proportion (95% CI)	λ Mean (Min, Max)
A.1	300	0.94 (0.88, 0.97)	0.017 (0, 0.13)	0.94 (0.88, 0.97)	0.026 (0, 0.10)
	500	0.97 (0.92, 0.99)	0.012 (0, 0.08)	0.96 (0.90, 0.98)	0.016 (0, 0.06)
	800	0.95 (0.89, 0.98)	0.004 (0, 0.04)	0.94 (0.88, 0.97)	0.010 (0, 0.04)
A.2	300	0.44 (0.35, 0.54)	0.033 (0, 0.20)	0.48 (0.38, 0.58)	0.034 (0, 0.17)
	500	0.67 (0.57, 0.75)	0.016 (0, 0.14)	0.92 (0.85, 0.96)	0.014 (0, 0.06)
	800	0.78 (0.69, 0.85)	0.010 (0, 0.09)	0.91 (0.84, 0.95)	0.012 (0, 0.04)
B.1	300	0.83 (0.74, 0.89)	0.043 (0, 0.41)	0.90 (0.83, 0.94)	0.083 (0, 0.50)
	500	0.77 (0.68, 0.84)	0.034 (0, 0.22)	0.90 (0.83, 0.94)	0.057 (0, 0.40)
	800	0.87 (0.79, 0.92)	0.029 (0, 0.17)	0.91 (0.84, 0.95)	0.041 (0, 0.19)
B.2	300	0.39 (0.30, 0.49)	0.052 (0, 0.35)	0.20 (0.13, 0.29)	0.085 (0, 0.40)
	500	0.33 (0.25, 0.43)	0.049 (0, 0.15)	0.37 (0.28, 0.47)	0.048 (0, 0.22)
	800	0.47 (0.38, 0.57)	0.031 (0, 0.16)	0.67 (0.57, 0.75)	0.031 (0, 0.33)
D	300	0.99 (0.95, 0.96)	0.004 (0, 0.15)	0.99 (0.95, 1.00)	0.002 (0, 0.12)
	500	1.00 (0.96, 1.00)	0.000 (0, 0.01)	1.00 (0.96, 1.00)	0.000 (0, 0.17)
	800	1.00 (0.96, 1.00)	0.000 (0, 0)	1.00 (0.96, 1.00)	0.000 (0, 0.02)

Table 3:

(Upper Panel) Probability (as %) of being assigned to active treatment using the ITR forest method. For each simulation scheme an ITR forest is trained on 300 observations and validated on 1000 observations with each forest consisting of 500 trees. Signal ratio of $\theta = 3$ (A and B schemes) or interaction effect $\beta = 3$ (C schemes) is used. IPWE probabilities are estimated for all schemes and AIPWE probabilities are shown in parentheses for EMR schemes B and C. (Lower Panel) Percent of all 1000 validation observations assigned to the correct treatment group based on majority voting for the ITR forest (top row) and an ITR tree (bottom row). In scheme C.1 the ITR tree using the AIPWE produced a null model. All subjects in simulation scheme C.1 should all benefit from treatment, a “True” column is therefore excluded. All simulation replicates include 10 excess noise variables.

ID	Scheme A.1		Scheme A.2		Scheme B.1		Scheme B.2		Scheme C.1		Scheme C.2	
	Pr ^I IPWE	True ²	Pr ^I IPWE	True ²	Pr ^I IPWE(AIPWE)	True ²	Pr ^I IPWE(AIPWE)	True ²	Pr ^I IPWE(AIPWE)	True ²	Pr ^I IPWE(AIPWE)	True ²
1	35.9	0	100	1	13.6 (19.6)	0	72.4 (72.7)	0	99.2 (100)	20.4 (16.0)	0	0
2	96.0	1	100	1	22.4 (20.6)	0	97.8 (100)	1	81.0 (94.2)	29.7 (32.0)	0	0
3	43.3	0	96.5	1	90.4 (89.1)	1	42.2 (37.4)	0	91.1 (92.5)	17.4 (15.4)	0	0
4	44.0	0	36.0	0	18.0 (21.0)	0	19.4 (17.6)	0	86.1 (97.5)	61.1 (54.0)	1	1
5	49.0	0	42.5	0	58.6 (63.5)	1	98.6 (99.6)	1	98.7 (100)	36.3 (38.1)	0	0
6	95.3	1	49.1	0	13.0 (20.8)	0	45.9 (42.1)	0	88.6 (79.2)	56.5 (46.9)	1	1
7	27.5	0	42.5	0	24.4 (26.0)	0	80.6 (84.5)	1	97.5 (97.5)	32.1 (37.5)	0	0
8	36.2	0	33.3	0	29.4 (31.0)	0	96.4 (97.8)	1	61.2 (70.0)	56.7 (60.2)	1	1
9	38.6	0	31.6	0	16.2 (23.8)	0	91.0 (96.8)	1	65.0 (82.5)	21.4 (20.1)	0	0
10	99.7	1	93.0	1	83.6 (80.6)	1	96.4 (95.3)	1	81.9 (97.5)	43.1 (41.0)	1	1
% Correct												
ITR Forest	99.7%		86.3%		100% (99.7%)		86.5% (89.5%)		94.3% (90.1%)		84.3% (87.5%)	
ITR Forest	100%		93.5%		100% (100%)		80.9% (93.6%)		95.0% (-)		75.8% (87.7%)	

ITR: individualized treatment rule; **IPWE**: inverse probability weighted estimator; **AIPWE**: augmented inverse probability weighted estimator.

I_{Pr} : % of trees voting for treatment among the 500 trees in ITR forest.

2_{True} : the known optimal treatment assignment. A treatment assignment of “1” indicates active treatment and “0” indicates control.

Table 4:

Average maximized value (standard deviation) for ITR forest, ITR tree, MIDAs by Laber and Zhao (2015), IT tree by Su et al. (2009), OWL via SVM by Zhao et al. (2012), and random guessing, over 100 simulation replicates. Models were trained on $n = 300$ observations with $\theta = 3$ (A and B schemes) or $\beta = 3$ (C schemes), along with 10, 25, or 50 excess noise predictors (p). Reported mean maximized value and standard derivation were estimated based on the validation samples of $n = 1000$ and 100 simulation replicates. Note that for all schemes larger values are desired. The IT tree method failed to identify a rule in over 50% of replicates for C.1 and D so results are not included.

Model	p	ITR Forest	ITR Tree		MIDAs	IT Tree	OWL	Random
			IPWE	AIPWE				
A.2	10	5.427 (0.06)	5.181 (0.09)	5.221 (0.08)	5.196 (0.09)	5.288 (0.04)	4.994 (0.06)	4.131 (0.06)
A.2	25	5.399 (0.06)	5.166 (0.08)	5.200 (0.07)	5.184 (0.08)	5.297 (0.02)	4.912 (0.06)	4.142 (0.06)
A.2	50	5.385 (0.08)	5.152 (0.12)	5.203 (0.09)	5.178 (0.09)	5.284 (0.06)	4.828 (0.06)	4.131 (0.07)
B.2	10	6.586 (0.05)	6.406 (0.05)	6.404 (0.05)	6.264 (0.10)	6.387 (0.06)	6.257 (0.07)	5.326 (0.09)
B.2	25	6.557 (0.05)	6.384 (0.05)	6.389 (0.04)	6.268 (0.12)	6.286 (0.05)	6.174 (0.10)	5.350 (0.09)
B.2	50	6.546 (0.05)	6.372 (0.08)	6.392 (0.05)	6.257 (0.09)	6.391 (0.05)	6.098 (0.09)	5.343 (0.09)
C.1	10	8.023 (0.09)	7.923 (0.13)	7.911 (0.12)	7.816 (0.10)	-	7.803 (0.10)	7.640 (0.05)
C.1	25	8.015 (0.07)	7.891 (0.13)	7.917 (0.12)	7.845 (0.11)	-	7.764 (0.08)	7.640 (0.07)
C.1	50	7.994 (0.09)	7.884 (0.15)	7.900 (0.15)	7.828 (0.14)	-	7.720 (0.09)	7.659 (0.07)
C.2	10	12.72 (0.24)	12.08 (0.42)	12.19 (0.36)	11.22 (0.44)	12.63 (0.27)	12.31 (0.31)	8.905 (0.23)
C.2	25	12.65 (0.31)	12.06 (0.31)	12.10 (0.35)	11.20 (0.51)	12.63 (0.23)	12.04 (0.30)	8.881 (0.22)
C.2	50	12.55 (0.30)	11.96 (0.35)	12.10 (0.39)	11.20 (0.40)	12.53 (0.25)	11.87 (0.28)	8.878 (0.21)
D	50	5.048 (0.08)	5.081 (0.09)	5.055 (0.08)	5.073 (0.05)	-	4.902 (0.08)	4.187 (0.14)

ITR: individualized treatment rule; **IPWE:** inverse probability weighted estimator; **AIPWE:** augmented inverse probability weighted estimator; **MIDAs:** minimum impurity decision assignments; **IT tree:** interaction tree; **OWL:** outcome weighted learning; **SVM:** support vector machine.

Table 5:

Probabilities of 10 held-out subjects being assigned to treatment group using a 500 tree ITR forest for DURABLE and EMR studies. Models were trained using all but 10 held-out subjects. Shown are:

DURABLE RCT				EMR			
Subject	Pr ¹	ITR Tree ²	Original ³	Subject	Pr ¹ IPWE(AIPWE)	ITR Tree ²	Original ³
1	61.6	0	1	1	100 (100)	1	0
2	90.0	1	1	2	100 (99.2)	1	0
3	74.6	1	1	3	98.7 (95.7)	1	0
4	48.6	0	1	4	100 (100)	1	1
5	87.4	1	1	5	68.9 (70.4)	1	1
6	69.0	0	0	6	98.3 (94.3)	1	1
7	76.5	1	1	7	100 (100)	1	1
8	83.3	1	1	8	100 (100)	1	0
9	71.0	0	1	9	97.3 (93.2)	1	1
10	67.1	1	1	10	91.6 (87.9)	1	0

¹“Pr”: % trees voting for treatment over 500 trees in ITR forest.

²“ITR Tree”: Treatment assignment from ITR tree.

³“Original”: original treatment assignments. IPWE estimates are shown with AIPWE estimates in parentheses for EMR. IPWE and APIWE in “ITR Tree” procedure generate the same assignments in EMR study (7th column). For DURABLE trial, treatment assignment of “1” indicates LM75/25 and “0” indicates Glargine. For EMR, treatment assignment of “1” indicates GLP-1 and “0” indicates basal insulin.

ITR: individualized treatment rule; **RCT**: randomized controlled trial; **EMR**: electronic medical records; **IPWE**: inverse probability weighted estimator; **AIPWE**: augmented inverse probability weighted estimator; **DURABLE**: Assessing the DURAbility of basal vs lispro mix 75/25 insulin efficacy trial; **LM75/25**: mix of 75% insulin lispro protamine suspension and 25% lispro (LM75/25); **GLP-1**: glucagon-like peptide-1 receptor agonists.

Table 6:

Computational time, in seconds, for each ITR method. EMR data from section 4 was used to train each model on a desktop computer with an Intel(R) Core i7-6700 CPU @ 3.40GHz with 16 GB of RAM. EMR data consisted of $n = 5177$ observations with 14 covariates (4 continuous, 10 categorical).

	ITR Forest	ITR Tree		MIDAs	IT Tree	OWL
		IPWE	AIPWE			
Time (seconds)	37.8	22.9	23.2	24.0	1.6	538 [*]

^{*} OWL model failed to converge after 538 seconds.

ITR: individualized treatment rule; **IPWE**: inverse probability weighted estimator; **AIPWE**: augmented inverse probability weighted estimator; **MIDAs**: minimum impurity decision assignments; **IT tree**: interaction tree; **OWL**: outcome weighted learning.