



A sparse additive model for treatment effect-modifier selection

HYUNG PARK*, EVA PETKOVA, THADDEUS TARPEY, R. TODD OGDEN

*Division of Biostatistics, Department of Population Health, New York University, New York, NY, USA and
Department of Biostatistics, Columbia University, New York, NY, USA*
parkh15@nyu.edu

SUMMARY

Sparse additive modeling is a class of effective methods for performing high-dimensional nonparametric regression. This article develops a sparse additive model focused on estimation of treatment effect modification with simultaneous treatment effect-modifier selection. We propose a version of the sparse additive model uniquely constrained to estimate the interaction effects between treatment and pretreatment covariates, while leaving the main effects of the pretreatment covariates unspecified. The proposed regression model can effectively identify treatment effect-modifiers that exhibit possibly nonlinear interactions with the treatment variable that are relevant for making optimal treatment decisions. A set of simulation experiments and an application to a dataset from a randomized clinical trial are presented to demonstrate the method.

Keywords: Individualized treatment rules; Sparse additive models; Treatment effect-modifiers; Biomarkers

1. INTRODUCTION

Identification of patient characteristics influencing treatment responses, which are often termed treatment effect-modifiers or treatment effect-moderators, is a top research priority in precision medicine. In this article, we develop a flexible yet simple and intuitive regression approach to identifying treatment effect-modifiers from a potentially large number of pretreatment patient characteristics. In particular, we utilize a sparse additive model (Ravikumar and others, 2009) to conduct effective treatment effect-modifier selection.

The clinical motivation behind the development of a high-dimensional additive regression model that can handle a large number of pretreatment covariates, specifically designed to model treatment effect modification and variable selection, is a randomized clinical trial (Trivedi and others, 2016) for treatment of major depressive disorder. A large number of baseline patient characteristics were collected from each participant prior to randomization and treatment allocation. The primary goal of this study, and thus the goal of our proposed method, is to discover biosignatures of heterogeneous treatment response, and to develop an individualized treatment rule (ITR; e.g., Qian and Murphy 2011;

*To whom correspondence should be addressed.

Zhao *and others* 2012; Laber and Zhao 2015; Kosorok and Laber 2019; Park *and others* 2020) for future patients based on biosignatures of the heterogeneous treatment response, a key aspect of precision medicine (e.g., Ashley 2015; Fernandes *and others* 2017, among others). In particular, discovering and identifying which pretreatment patient characteristics influence treatment effects have the potential to significantly enhance clinical reasoning in practice (see, e.g., Royston and Sauerbrei, 2008).

The major challenge in efficiently modeling treatment effect modification from a clinical trial study is that variability due to treatment effect modification (i.e., the treatment-by-covariates interaction effects on outcomes), that is essential for developing ITRs (Qian and Murphy, 2011), is typically dwarfed by a relatively large degree of non-treatment-related variability (i.e., the main effects of the pretreatment covariates on treatment outcomes). In particular, in regression, due to potential confounding between the main effect of the covariates and the treatment-by-covariates interaction effect, misspecification of the covariate main effect may significantly influence estimation of the treatment-by-covariates interaction effect.

A simple and elegant linear model-based approach to modeling the treatment-by-covariates interactions, termed the modified covariate (MC) method, that is robust against model misspecification of the covariates' main effects was developed by Tian *and others* (2014). The method utilizes a simple parameterization and bypasses the need to model the main effects of the covariates. See also Murphy (2003); Lu *and others* (2011), Shi *and others* (2016), and Jeng *and others* (2018) for the similar linear model-based approaches to estimating the treatment-by-covariates interactions. However, these approaches assume a stringent linear model to specify the treatment-by-covariates interaction effects and are limited to the binary-valued treatment variable case only. In this work, we extend the optimization framework of Tian *and others* (2014) for modeling the treatment-by-covariates interactions using a more flexible regression setting based on additive models (Hastie and Tibshirani, 1999) that can accommodate more than two treatment conditions, while allowing an unspecified main effect of the covariates. Additionally, via an appropriate regularization, the proposed approach simultaneously achieves treatment effect-modifier selection in estimation.

In Section 2, we introduce an additive model that has both the unspecified main effect of the covariates and the treatment-by-covariates interaction effect additive components. Then, we develop an optimization framework specifically targeting the interaction effect additive components of the model, with a sparsity-inducing regularization parameter to encourage sparsity in the set of component functions. In Section 3, we develop a coordinate decent algorithm to estimate the interaction effect part of the model and consider an optimization strategy for ITRs. In Section 4, we illustrate the performance of the method in terms of treatment effect-modifier selection and estimation of ITRs through simulation examples. Section 5 provides an illustrative application from a depression clinical trial, and the article concludes with discussion in Section 6.

2. MODELS

Let $A \in \{1, \dots, L\}$ denote a treatment variable assigned with associated probabilities (π_1, \dots, π_L) , $\sum_{a=1}^L \pi_a = 1$, and $\pi_a > 0$, and let $\mathbf{X} = (X_1, \dots, X_p)^\top \in \mathbb{R}^p$ denote pretreatment covariates, independent of A (as in the case of a randomized trial). If the treatment assignment depends on the covariates, then the probabilities $\{\pi_a\}_{a \in \{1, \dots, L\}}$ can be replaced by propensity scores $\{\pi_a(\mathbf{X})\}_{a \in \{1, \dots, L\}}$ which would typically need to be estimated from data. We let $Y^{(a)} \in \mathbb{R}$ ($a = 1, \dots, L$) be the potential outcome if the patient received treatment $A = a$ ($a = 1, \dots, L$); we only observe $Y = Y^{(A)}$, A and \mathbf{X} . Throughout the article, without loss of generality, we assume that a larger value of the treatment outcome Y is desirable, and that $\mathbb{E}[Y|A = a] = 0$ ($a = 1, \dots, L$), i.e., the main effect of A on the outcome is centered at 0; this is only to suppress the treatment a -specific intercepts in regression models in order to simplify the exposition, and

can be achieved by removing the treatment level a -specific means from $Y \in \mathbb{R}$. We model the treatment outcome Y by the following additive model:

$$\mathbb{E}[Y|\mathbf{X}, A = a] = \underbrace{\mu^*(\mathbf{X})}_{X \text{ "main" effect}} + \underbrace{\sum_{j=1}^p g_{j,a}^*(X_j)}_{A\text{-by-}X \text{ interaction effect}} \quad (a = 1, \dots, L). \quad (2.1)$$

In model (2.1), the first term $\mu^*(\mathbf{X})$, which in fact will not need to be specified in our exposition, does not depend on the treatment variable A and thus the A -by- X interaction effects are determined only by the second component $\sum_{j=1}^p g_{j,a}^*(X_j)$. In terms of modeling treatment effect modification, the term $\sum_{j=1}^p g_{j,a}^*(X_j)$ in (2.1) corresponds to the “signal” component, whereas the term $\mu^*(\mathbf{X})$ corresponds to a “nuisance” component. In particular, the optimal ITR, under model (2.1), can be shown to satisfy $\mathcal{D}^{\text{opt}}(\mathbf{X}) = \arg \max_{a \in \{1, \dots, L\}} \sum_{j=1}^p g_{j,a}^*(X_j)$, which does not involve the term $\mu^*(\mathbf{X})$. Therefore, the A -by- X interaction effect term $\sum_{j=1}^p g_{j,a}^*(X_j)$ shall be the primary estimation target of model (2.1).

In (2.1), for each individual covariate X_j , we utilize a treatment a -specific smooth $g_{j,a}^*$ separately for each treatment condition $a \in \{1, \dots, L\}$. However, it is useful to treat the set of treatment-specific smooths for X_j as a single unit, i.e., a single component function $g_j^* = \{g_{j,a}^*\}_{a \in \{1, \dots, L\}}$, for the purpose of treatment effect-modifier variable selection.

In model (2.1), to separate $\mu^*(\mathbf{X})$ from the component $\sum_{j=1}^p g_{j,a}^*(X_j)$ and obtain an identifiable representation, without loss of generality, we assume that the set of treatment a -specific smooths $\{g_{j,a}^*\}_{a \in \{1, \dots, L\}}$ of the j th component function $g_{j,a}^*(X_j)$ satisfies a condition (almost surely):

$$\mathbb{E}[g_{j,a}^*(X_j)|X_j] = \sum_{a=1}^L \pi_a g_{j,a}^*(X_j) = 0 \quad (j = 1, \dots, p). \quad (2.2)$$

Condition (2.2) implies $\mathbb{E}[\sum_{j=1}^p g_{j,a}^*(X_j)|\mathbf{X}] = 0$ and separates the A -by- X interaction effect component, $\sum_{j=1}^p g_{j,a}^*(X_j)$, from the X “main” effect component, $\mu^*(\mathbf{X})$, in model (2.1). For model (2.1), we assume an additive noise structure $Y = \mathbb{E}[Y|\mathbf{X}, A] + \epsilon$, where ϵ is a zero-mean noise with a finite second moment.

Notation: For a single component j and general component functions $g_j = \{g_{j,a}\}_{a \in \{1, \dots, L\}}$, we define the L^2 norm of g_j as $\|g_j\| = \sqrt{\mathbb{E}[g_{j,A}^2(X_j)]}$, where expectation is taken with respect to the joint distribution of (A, X_j) . For a set of random variables (A, X_j) , let $\mathcal{H}_j = \{g_j \mid \mathbb{E}[g_{j,A}(X_j)] = 0, \|g_j\| < \infty\}$ with inner product on the space defined as $\langle g_j, f_j \rangle = \mathbb{E}[g_{j,A}(X_j)f_{j,A}(X_j)]$. Sometimes we also write $g_j := g_{j,A}(X_j)$ for the notational simplicity.

Under model (2.1) subject to (2.2), the component functions $\{g_j^*, j = 1, \dots, p\}$ associated with the A -by- X interaction effect can be viewed as the solution to the constrained optimization:

$$\begin{aligned} \{g_j^*\} &= \underset{\{g_j \in \mathcal{H}_j\}}{\operatorname{argmin}} \quad E \left[\left\{ Y - \mu^*(\mathbf{X}) - \sum_{j=1}^p g_{j,A}(X_j) \right\}^2 \right] \\ &\text{subject to} \quad \mathbb{E}[g_{j,A}(X_j)|X_j] = 0 \quad (j = 1, \dots, p), \end{aligned} \quad (2.3)$$

where $\mu^*(\mathbf{X})$ is given from the assumed model (2.1) (and is considered as fixed in (2.3)). Since the minimization in (2.3) is in terms of $\{g_j \in \mathcal{H}_j, j = 1, \dots, p\}$, the objective function part of the right-hand

side of (2.3) can be reduced to:

$$\begin{aligned} & \arg \min_{\{g_j \in \mathcal{H}_j\}} \mathbb{E} \left[Y^2 + \left\{ \sum_{j=1}^p g_{j,A}(X_j) \right\}^2 - 2 \left\{ \sum_{j=1}^p g_{j,A}(X_j) \right\} Y + 2 \left\{ \sum_{j=1}^p g_{j,A}(X_j) \right\} \mu^*(\mathbf{X}) \right] \\ &= \arg \min_{\{g_j \in \mathcal{H}_j\}} \mathbb{E} \left[Y^2 + \left\{ \sum_{j=1}^p g_{j,A}(X_j) \right\}^2 - 2 \left\{ \sum_{j=1}^p g_{j,A}(X_j) \right\} Y + 2 \mu^*(\mathbf{X}) \mathbb{E} \left[\sum_{j=1}^p g_{j,A}(X_j) | \mathbf{X} \right] \right] \\ &= \arg \min_{\{g_j \in \mathcal{H}_j\}} \mathbb{E} \left[Y^2 + \left\{ \sum_{j=1}^p g_{j,A}(X_j) \right\}^2 - 2 \left\{ \sum_{j=1}^p g_{j,A}(X_j) \right\} Y \right], \end{aligned}$$

where the second line is from an application of the iterated expectation rule conditioning on \mathbf{X} and the third line follows from the constraint $\mathbb{E}[g_{j,A}(X_j)|X_j] = 0$ ($j = 1, \dots, p$) on the right-hand side of (2.3). Therefore, representation (2.3) can be simplified to:

$$\begin{aligned} \{g_j^*\} &= \underset{\{g_j \in \mathcal{H}_j\}}{\operatorname{argmin}} \quad \mathbb{E} \left[\left\{ Y - \sum_{j=1}^p g_{j,A}(X_j) \right\}^2 \right] \\ &\text{subject to} \quad \mathbb{E}[g_{j,A}(X_j)|X_j] = 0 \quad (j = 1, \dots, p), \end{aligned} \tag{2.4}$$

Representation (2.4) of the component functions $\{g_j^*, j = 1, \dots, p\}$ of the underlying model (2.1) is particularly useful when the (high-dimensional) “nuisance” function μ^* in (2.1) is complicated and prone to specification error. Note,

$$\mathbb{E} \left[\mu^*(\mathbf{X}) \sum_{j=1}^p g_{j,A}(X_j) \right] = \mathbb{E} \left[\mu^*(\mathbf{X}) \sum_{j=1}^p \mathbb{E}[g_{j,A}(X_j)|\mathbf{X}] \right] = 0, \tag{2.5}$$

indicating that the component $\sum_{j=1}^p g_{j,A}(X_j) \in \mathcal{H}_1 + \dots + \mathcal{H}_p$ (subject to the identifiability constraint $\mathbb{E}[g_{j,A}(X_j)|X_j] = 0$) is structured to be orthogonal to the \mathbf{X} main effect $\mu^*(\mathbf{X})$. This orthogonality property is useful for estimating the additive model A -by- \mathbf{X} interaction effect $\sum_{j=1}^p g_{j,A}^*(X_j)$, in the presence of the unspecified $\mu^*(\mathbf{X})$ of model (2.1).

Under model (2.1), the potential treatment effect-modifiers among $\{X_j, j = 1, \dots, p\}$ enter the model only through the interaction effect term $\sum_{j=1}^p g_{j,A}^*(X_j)$ that associates the treatment A to the treatment outcome Y . Ravikumar and others (2009) proposed sparse additive modeling (SAM) for component selection in high-dimensional additive models with a large p . As in SAM, to deal with a large p and to achieve treatment effect-modifier selection, we impose sparsity on the set of component functions $\{g_j^*, j = 1, \dots, p\}$ associated with the interaction effect term of model (2.1), under the often practical and reasonable assumption that most covariates are irrelevant as treatment effect-modifiers. This sparsity structure on the index set $\{j\}$ for the nonzero component functions $\{g_j^* : g_j^* \neq 0\}$ can be usefully incorporated into the optimization-based criterion (2.4) in representing $\{g_j^*\}$:

$$\begin{aligned} \{g_j^*\} &= \underset{\{g_j \in \mathcal{H}_j\}}{\operatorname{argmin}} \quad \mathbb{E} \left[\left\{ Y - \sum_{j=1}^p g_{j,A}(X_j) \right\}^2 \right] + \lambda \sum_{j=1}^p \|g_j\| \\ &\text{subject to} \quad \mathbb{E}[g_{j,A}(X_j)|X_j] = 0 \quad (j = 1, \dots, p), \end{aligned} \tag{2.6}$$

for a sparsity-inducing parameter $\lambda \geq 0$. The term $\sum_{j=1}^p \|g_j\|$ in (2.6) behaves like an L^1 ball across different components $\{g_j, j = 1, \dots, p\}$ to encourage sparsity in the set of component functions. For example, a

large value of λ on the right-hand side of (2.6) will generate a sparse solution with many component functions g_j^* on the left-hand side set exactly to zero.

3. ESTIMATION

3.1. Model estimation

For each j , the minimizer g_j^* of the optimization problem (2.6) has a component-wise closed-form expression.

THEOREM 1 Given $\lambda \geq 0$, the minimizer $g_j^* \in \mathcal{H}_j$ of (2.6) satisfies (almost surely):

$$g_{j,A}^*(X_j) = \left[1 - \frac{\lambda}{\|f_j\|} \right]_+ f_{j,A}(X_j), \quad (3.7)$$

where

$$f_{j,A}(X_j) = \mathbb{E}[R_j|X_j, A] - \mathbb{E}[R_j|X_j], \quad (3.8)$$

in which

$$R_j = Y - \sum_{j' \neq j} g_{j',A}^*(X_{j'}) \quad (3.9)$$

represents the j th partial residual. In (3.7), $Z_+ = \max(0, Z)$ represents the positive part of Z .

Note that the $f_{j,A}(X_j)$ correspond to the projections of R_j onto \mathcal{H}_j subject to the constraint in (2.6). The proof of Theorem 1 is in the [Supplementary Materials](#) available at *Biostatistics* online.

The component-wise expression (3.7) for g_j^* suggests that we can employ a coordinate descent algorithm (e.g., [Tseng, 2001](#)) to solve (2.6). Given a sparsity parameter $\lambda \geq 0$, we can use a standard backfitting algorithm used in fitting additive models ([Hastie and Tibshirani, 1999](#)) that fixes the set of current approximates for $g_{j'}^*$ at all $j' \neq j$, and obtains a new approximate of g_j^* by equation (3.7), and iterates through all j until convergence. A sample version of the algorithm can be obtained by inserting sample estimates into the population expressions (3.9), (3.8), and (3.7) for each coordinate j , which we briefly describe next.

Given data (X_{ij}, A_i) ($i = 1, \dots, n; j = 1, \dots, p$), for each j , let $\hat{R}_j = Y - \sum_{j' \neq j} \hat{g}_{j',A}^*(X_{j'})$, corresponding to the data-version of the j th partial residual R_j in (3.9), where $\hat{g}_{j'}^*$ represents a current estimate for $g_{j'}^*$. We estimate g_j^* in (3.7) in two steps: (i) estimate the function f_j in (3.8); (ii) plug the estimate of f_j into $\left[1 - \frac{\lambda}{\|f_j\|} \right]_+$ in (3.7), to obtain the soft-thresholded estimate \hat{g}_j^* .

Although any linear smoothers can be utilized to obtain the estimators $\{\hat{g}_j^*\}$ as described in Remark 1 at the end of this section, in this article, we shall focus on regression spline-type estimators which are particularly simple and computationally efficient to implement. Specifically, for each j , the function $g_j \in \mathcal{H}_j$ on the right-hand side of (2.6) will be represented by:

$$g_{j,a}(X_j) = \Psi_j(X_j)^\top \theta_{j,a} \quad (a = 1, \dots, L) \quad (3.10)$$

for some prespecified d_j -dimensional basis $\Psi_j(\cdot) \in \mathbb{R}^{d_j}$ (e.g., B -spline basis on evenly spaced knots on a bounded range for X_j) and a set of unknown treatment a -specific basis coefficients $\{\theta_{j,a} \in \mathbb{R}^{d_j}\}_{a \in \{1, \dots, L\}}$.

Given representation (3.10) for the component function g_j , the constraint $\mathbb{E}[g_{j,A}(X_j)|X_j] = 0$ in (2.6) can be simplified to $\mathbb{E}[\boldsymbol{\theta}_{j,A}] = \sum_{a=1}^L \pi_a \boldsymbol{\theta}_{j,a} = \mathbf{0}$. This constraint can be written succinctly in a matrix form as

$$\boldsymbol{\pi}^{(j)} \boldsymbol{\theta}_j = \mathbf{0}, \quad (3.11)$$

where $\boldsymbol{\theta}_j := (\boldsymbol{\theta}_{j,1}^\top, \boldsymbol{\theta}_{j,2}^\top, \dots, \boldsymbol{\theta}_{j,L}^\top)^\top \in \mathbb{R}^{d_j L}$ is the vectorized version of the basis coefficients $\{\boldsymbol{\theta}_{j,a}\}_{a \in \{1, \dots, L\}}$ in (3.10), and the $d_j \times d_j L$ matrix $\boldsymbol{\pi}^{(j)} := (\pi_1 \mathbf{I}_{d_j}; \pi_2 \mathbf{I}_{d_j}; \dots; \pi_L \mathbf{I}_{d_j})$ where \mathbf{I}_{d_j} denotes the $d_j \times d_j$ identity matrix.

Let the $n \times d_j$ matrices $\mathbf{D}_{j,a}$ ($a = 1, \dots, L$) denote the evaluation matrices of the basis function $\boldsymbol{\Psi}_j(\cdot)$ on X_{ij} ($i = 1, \dots, n$) specific to the treatment $A = a$ ($a = 1, \dots, L$), whose i th row is the $1 \times d_j$ vector $\boldsymbol{\Psi}_j(X_{ij})^\top$ if $A_i = a$, and a row of zeros $\mathbf{0}^\top$ if $A_i \neq a$. Then the column-wise concatenation of the design matrices $\{\mathbf{D}_{j,a}\}_{a \in \{1, \dots, L\}}$, i.e., the $n \times d_j L$ matrix $\mathbf{D}_j = (\mathbf{D}_{j,1}; \mathbf{D}_{j,2}; \dots; \mathbf{D}_{j,L})$, defines the model matrix associated with the vectorized model coefficient $\boldsymbol{\theta}_j \in \mathbb{R}^{d_j L}$, vectorized across $\{\boldsymbol{\theta}_{j,a}\}_{a \in \{1, \dots, L\}}$ in representation (3.10). Then, we can represent the function $g_{j,A}(X_j)$ in (3.10), based on the sample data, by the length- n vector:

$$\mathbf{g}_j = \mathbf{D}_j \boldsymbol{\theta}_j \quad (3.12)$$

subject to the linear constraint (3.11).

The linear constraint (3.11) on $\boldsymbol{\theta}_j$ can be conveniently absorbed into the model matrix \mathbf{D}_j in (3.12) by reparametrization, as we describe next. We can find a $d_j L \times d_j(L-1)$ basis matrix $\mathbf{Z}^{(j)}$, such that if we set $\boldsymbol{\theta}_j = \mathbf{Z}^{(j)} \tilde{\boldsymbol{\theta}}_j$ for any arbitrary vector $\tilde{\boldsymbol{\theta}}_j \in \mathbb{R}^{d_j(L-1)}$, then the vector $\boldsymbol{\theta}_j \in \mathbb{R}^{d_j L}$ automatically satisfies the constraint (3.11) $\boldsymbol{\pi}^{(j)} \boldsymbol{\theta}_j = \mathbf{0}$. Such a basis matrix $\mathbf{Z}^{(j)}$ can be constructed by a QR decomposition of the matrix $\boldsymbol{\pi}^{(j)\top}$. Then representation (3.12) can be reparametrized, in terms of the unconstrained vector $\tilde{\boldsymbol{\theta}}_j \in \mathbb{R}^{d_j(L-1)}$, by replacing \mathbf{D}_j in (3.12) with the reparametrized model matrix $\tilde{\mathbf{D}}_j = \mathbf{D}_j \mathbf{Z}^{(j)}$:

$$\mathbf{g}_j = \tilde{\mathbf{D}}_j \tilde{\boldsymbol{\theta}}_j. \quad (3.13)$$

Theorem 1 and Section 2 of Supplementary Materials available at *Biostatistics* online indicate that the coordinate-wise minimizer \mathbf{g}_j^* of (2.6) can be estimated based on the sample by

$$\hat{\mathbf{g}}_j^* = \left[1 - \frac{\lambda}{\sqrt{\frac{1}{n} \|\hat{\mathbf{f}}_j\|^2}} \right]_+ \hat{\mathbf{f}}_j, \quad (3.14)$$

where

$$\hat{\mathbf{f}}_j = \tilde{\mathbf{D}}_j (\tilde{\mathbf{D}}_j^\top \tilde{\mathbf{D}}_j)^{-1} \tilde{\mathbf{D}}_j^\top \hat{\mathbf{R}}_j \quad (3.15)$$

in which $\hat{\mathbf{R}}_j = \mathbf{Y} - \sum_{j' \neq j} \hat{\mathbf{g}}_{j'}^*$ is the estimated j th partial residual vector. In (3.14), the norm $\|\hat{f}_j\|$ of (3.7) is estimated by the vector norm $\sqrt{\frac{1}{n} \|\hat{\mathbf{f}}_j\|^2}$, and the shrinkage factor $(\hat{s}_j^{(\lambda)} =) \left[1 - \frac{\lambda}{\|\hat{f}_j\|} \right]_+$ of (3.7) is estimated by $(\hat{s}_j^{(\lambda)} =) \left[1 - \frac{\lambda}{\sqrt{\frac{1}{n} \|\hat{\mathbf{f}}_j\|^2}} \right]_+$.

Based on the sample counterpart (3.14) of the coordinate-wise solution (3.7), a highly efficient coordinate descent algorithm can be conducted to simultaneously estimate all the component functions

$\{\mathbf{g}_j^*, j = 1, \dots, p\}$ in (2.6). At convergence of the coordinate descent, we have a basis coefficient estimate associated with the representation (3.13),

$$\hat{\boldsymbol{\theta}}_j = \hat{s}_j^{(\lambda)} (\tilde{\mathbf{D}}_j^\top \tilde{\mathbf{D}}_j)^{-1} \tilde{\mathbf{D}}_j^\top \hat{\mathbf{R}}_j \quad (3.16)$$

which in turn implies an estimate

$$\hat{\boldsymbol{\theta}}_j = (\hat{\boldsymbol{\theta}}_{j,1}^\top, \hat{\boldsymbol{\theta}}_{j,2}^\top, \dots, \hat{\boldsymbol{\theta}}_{j,L}^\top)^\top = \mathbf{Z}^{(j)\top} \hat{\boldsymbol{\theta}}_j$$

for the basis coefficient associated with the representation (3.12). This gives an estimate of the treatment a -specific function $g_{j,a}^*(\cdot)$ ($a = 1, \dots, L$) in model (2.1):

$$\hat{g}_{j,a}^*(\cdot) = \boldsymbol{\Psi}_j(\cdot)^\top \hat{\boldsymbol{\theta}}_{j,a} \quad (a = 1, \dots, L) \quad (3.17)$$

estimated within the class of functions (3.10) for a given tuning parameter λ , which controls the shrinkage factor $\hat{s}_j^{(\lambda)}$ in (3.16). We summarize the computational procedure for the coordinate descent in Algorithm 1.

Algorithm 1 Coordinate descent

- 1: **Input:** Data $\mathbf{X} \in \mathbb{R}^n \times \mathbb{R}^p$, $\mathbf{A} \in \mathbb{R}^n$, $\mathbf{Y} \in \mathbb{R}^n$, and tuning parameter λ .
 - 2: **Output:** Fitted functions $\{\hat{\mathbf{g}}_j^*, j = 1, \dots, p\}$.
 - 3: Initialize $\hat{\mathbf{g}}_j^* = \mathbf{0} \forall j$; pre-compute the smoother matrices $\tilde{\mathbf{D}}_j (\tilde{\mathbf{D}}_j^\top \tilde{\mathbf{D}}_j)^{-1} \tilde{\mathbf{D}}_j^\top$ in (3.15) $\forall j$.
 - 4: **while** until convergence of $\{\hat{\mathbf{g}}_j^*, j = 1, \dots, p\}$, **do** iterate through $j = 1, \dots, p$:
 - 5: Compute the partial residual $\hat{\mathbf{R}}_j = \mathbf{Y} - \sum_{j' \neq j} \hat{\mathbf{g}}_{j'}^*$
 - 6: Compute $\hat{\mathbf{f}}_j$ in (3.15); then compute the thresholded estimate $\hat{\mathbf{g}}_j^*$ in (3.14).
-

In Algorithm 1, the projection matrices $\tilde{\mathbf{D}}_j (\tilde{\mathbf{D}}_j^\top \tilde{\mathbf{D}}_j)^{-1} \tilde{\mathbf{D}}_j^\top$ ($j = 1, \dots, p$) only need to be computed once and therefore the coordinate descent can be performed efficiently. In (3.14), if the shrinkage factor $\hat{s}_j^{(\lambda)} = \left[1 - \frac{\lambda \sqrt{n}}{\|\tilde{\mathbf{f}}_j\|}\right]_+ = 0$, the associated j th covariate is absent from the model. The tuning parameter $\lambda \geq 0$ for treatment effect-modifier selection can be chosen to minimize an estimate of the expected squared error of the fitted models, $\mathbb{E}[\{Y - \sum_{j=1}^p \hat{g}_{j,A}^*(X_j)\}^2]$, over a dense grid of λ 's, estimated, for example, by cross-validation. Alternatively, one can utilize the network information criterion (Murata and Amari, 1994) which is a generalization of the Akaike information criterion in approximating the prediction error, for the case where the true underlying model, i.e., model (2.1), is not necessarily in the class of candidate models. Throughout the article, λ is chosen to minimize 10-fold cross-validated prediction error of the fitted models.

REMARK 1 For coordinate descent, any linear smoothers can be utilized to obtain the sample counterpart (3.14) of the coordinate-wise solution (3.7), i.e., the method is not restricted to regression splines. To estimate the function f_j in (3.8), we can estimate the first term $\mathbb{E}[R_j | X_j, A = a]$ in (3.8), using a 1D nonparametric smoother for each treatment level $a \in \{1, \dots, L\}$ separately, based on the data (\hat{R}_{ij}, X_{ij}) ($i \in \{i : A_i = a\}$) corresponding to the data from the a th treatment condition; we can also estimate the second term $-\mathbb{E}[R_j | X_j]$ in (3.8) based on the data $(\hat{R}_{ij}, X_{ij})_{i \in \{1, \dots, n\}}$ which corresponds to the set of data from all treatment conditions, using a 1D nonparametric smoother. Adding these two estimates evaluated at the n observed values of (X_{ij}, A_i) ($i = 1, \dots, n$) gives an estimate $\hat{\mathbf{f}}_j$ in (3.14). Then, we can compute the associated estimate $\hat{\mathbf{g}}_j^*$, which allows implementation of the coordinate descent in Algorithm 1.

3.2. Individualized treatment rule estimation

For a single time decision point, an ITR, which we denote by $\mathcal{D}(\mathbf{X}) : \mathbb{R}^p \mapsto \{1, \dots, L\}$, maps an individual with pretreatment characteristics \mathbf{X} to one of the L available treatment options. One natural measure for the effectiveness of an ITR \mathcal{D} in precision medicine is the so-called “value” (V) function (Murphy, 2005):

$$V(\mathcal{D}) = \mathbb{E}[\mathbb{E}[Y|\mathbf{X}, A = \mathcal{D}(\mathbf{X})]], \quad (3.18)$$

which is the expected treatment response under a given ITR \mathcal{D} . The optimal ITR \mathcal{D} , which we write as \mathcal{D}^{opt} , can be naturally defined to be the rule that maximizes the value $V(\mathcal{D})$ (3.18). Such an optimal rule \mathcal{D}^{opt} satisfies:

$$\mathcal{D}^{\text{opt}}(\mathbf{X}) = \arg \max_{a \in \{1, \dots, L\}} \mathbb{E}[Y|\mathbf{X}, A = a]. \quad (3.19)$$

Much work has been carried out to develop methods for estimating the optimal ITR (3.19) using data from randomized clinical trials. Machine learning approaches to estimating (3.19) are often framed in the context of a (weighted) classification problem (Zhang and others, 2012; Zhao and others, 2019), where the function $\mathcal{D}^{\text{opt}}(\mathbf{X})$ in (3.19) is regarded as the optimal classification rule given \mathbf{X} for the treatment assignment with respect to the objective function (3.18). These classification-based approaches to optimizing ITRs include the outcome-weighted learning (OWL) (e.g., Zhao and others, 2012, 2015; Song and others, 2015; Liu and others, 2018) based on support vector machines (SVMs), tree-based classification (e.g., Laber and Zhao, 2015), and adaptive boosting (Kang and others, 2014), among others.

Under model (2.1), $\mathcal{D}^{\text{opt}}(\mathbf{X})$ in (3.19) is: $\mathcal{D}^{\text{opt}}(\mathbf{X}) = \arg \max_{a \in \{1, \dots, L\}} \sum_{j=1}^p g_{j,a}^*(X_j)$, which can be estimated by: $\hat{\mathcal{D}}^{\text{opt}}(\mathbf{X}) = \arg \max_{a \in \{1, \dots, L\}} \sum_{j=1}^p \hat{g}_{j,a}^*(X_j)$, where $\hat{g}_{j,a}^*(\cdot)$ is given in (3.17) at the convergence of the Algorithm 1. This estimator can be viewed as a regression-based approach to estimating (3.19), that approximates the conditional expectations $\mathbb{E}[Y|\mathbf{X}, A = a]$ ($a = 1, \dots, L$) based on the additive model (2.1), while maintaining robustness with respect to model misspecification of the “nuisance” function μ^* in (2.1) via representation (2.6) for the “signal” components g_j^* . We illustrate the performance of this ITR estimator $\hat{\mathcal{D}}^{\text{opt}}$ with respect to the value function (3.18), through a set of simulation studies in Section 4.2.

3.3. Feature selection and transformation for individualized treatment rules

Although machine learning approaches that attempt to directly maximize (3.18) without assuming some specific structure on $\mathbb{E}[Y|\mathbf{X}, A = a]$ ($a = 1, \dots, L$) (unlike most of the regression-based approaches) are highly appealing, common machine learning approaches used in optimizing ITRs, including SVMs utilized in the OWL, are often hard to scale to large datasets, due to their taxing computational time. In particular, SVMs are viewed as “shallow” approaches (as opposed to a “deep” learning method that utilizes a learning model with many representational layers) and successful applications of SVMs often require first extracting useful representations for their input data manually or through some data-driven feature transformation (a step called feature engineering) (see, e.g., Kuhn and Johnson, 2019) to have more discriminatory power. Generally, selection and transformation of relevant features can increase the performance, scale, and speed of a machine learning procedure.

As an added value, the proposed regression (2.6) based on model (2.1) provides a practical feature selection and transformation learning technique for optimizing ITRs. The set of component functions $\{g_j^*, j = 1, \dots, p\}$ in (2.6) can be used to define data-driven feature transformation functions for the original features $\{X_j, j = 1, \dots, p\}$. The resulting transformed features can be used as inputs to a machine learning algorithm for optimizing ITRs and can lead to good results in practical situations.

In particular, we note that for each j , the component function g_j^* in (2.6) is defined separately from the X_j main effect function μ^* in (2.1). Therefore, the corresponding transformed feature variable $g_{j,1}^*(X_j)$, which represents the j th feature X_j in the new space, highlights only the “signal” nonlinear effect of X_j associated with the A -by- X_j interactions on their effects on the outcome that is relevant to estimating \mathcal{D}^{opt} , while excluding the X_j main effect that is irrelevant to the ITR development. This “de-noising” procedure for each variable X_j can be very appealing, since irrelevant or partially relevant features can negatively impact the performance of a machine learning algorithm. Moreover, a relatively large value of the tuning parameter $\lambda > 0$ in (2.6) would imply a set of sparse component functions $\{g_j^*\}$, providing a means of feature selection for ITRs. For the most common case of $L = 2$ (binary treatment), we have $g_{j,2}^*(X_j) = -\pi_2^{-1}\pi_1 g_{j,1}^*(X_j)$ implied by the constraint (2.2) that we impose, which is simply a scalar-scaling of the function $g_{j,1}^*(X_j)$; this implies that, for each j , the mapping $X_j \mapsto g_{j,1}^*(X_j)$ specifies the feature transformation of X_j . We demonstrate the utility of this feature selection/transformation, which we use as an input to the OWL approach to optimizing ITRs, through a set of simulation studies in Section 4.2 and a real data application in Section 5.

4. SIMULATION STUDY

4.1. Treatment effect-modifier selection performance

In this section, we will report simulation results illustrating the performance of the treatment effect-modifier selection. The complexity of the model for studying the A -by- \mathbf{X} interactions can be summarized in terms of the size of the index set for the component functions $\{g_j^*, j = 1, \dots, p\}$ that are not identically zero. We can ascertain the performance of a treatment effect-modifier selection method in terms of these component functions correctly or incorrectly estimated as nonzero. To generate the data, we use the following model:

$$Y = \sum_{j=1}^{10} \cos(X_j) + g_{1,A}^*(X_1) + g_{2,A}^*(X_2) + g_{3,A}^*(X_3) + \epsilon \quad A \in \{1, 2\}, \quad (4.20)$$

where X_j ($j = 1, \dots, p$), $p \in \{50, 200\}$ are generated from independent $\text{Unif}[-\pi/2, \pi/2]$, and the treatment variable $A \in \{1, 2\}$ is generated independently of \mathbf{X} and the error term $\epsilon \sim \mathcal{N}(0, 0.5^2)$, such that $\Pr(A = 1) = \Pr(A = 2) = 1/2$. We set $g_{1,A}^*(X_1) = (A - 1.5)X_1$, $g_{2,A}^*(X_2) = (A - 1.5) \{I_{(X_2 \leq 1.3)} 0.05e^{(X_2 - 1.3)} + I_{(X_2 > 1.3)} e^{4(X_2 - 1.3)} - 1\}$ and $g_{3,A}^*(X_3) = (A - 1.5) \{2e^{-X_3^2} - 1\}$. The graphs of these functions are displayed in Figure S.1 of Supplementary Materials available at *Biostatistics* online.

Under model (4.20), there are only three true treatment effect-modifiers, X_1, X_2 , and X_3 . The other $p - 3$ covariates are “noise” covariates, that are not consequential for optimizing ITRs. Also, in (4.20), there are 10 covariates X_j ($j = 1, \dots, 10$), among the p covariates, associated with the \mathbf{X} main effects. Under the setting (4.20), the contribution to the variance of the outcome from the \mathbf{X} main effect component was about 2.3 times greater than that from the A -by- \mathbf{X} interaction effect component.

We consider two approaches to treatment effect-modifier selection: (i) the proposed additive regression approach (2.6) that specifies a sparse set of functions $\{g_j^*, j = 1, \dots, p\}$, estimated via Algorithm 1, with the dimension of the cubic B -spline basis Ψ_j in (3.10) set to be $d_j = 6$ ($j = 1, \dots, p$); and (ii) the linear regression (MC) approach of Tian and others (2014),

$$\underset{\{\beta_j \in \mathbb{R}\}}{\text{minimize}} \quad \mathbb{E} \left[\left\{ Y - \sum_{j=1}^p (A - 1.5) X_j \beta_j \right\}^2 \right] + \lambda \sum_{j=1}^p |\beta_j|, \quad (4.21)$$




Fig. 1. The proportion of the three relevant covariates (i.e., X_1, X_2 , and X_3) correctly selected (the “true positives”; the two gray panels), and the $p - 3$ irrelevant covariates (i.e., X_4, X_5, \dots, X_p) incorrectly selected (the “false positives”; the two white panels), respectively (and ± 1 standard deviation), as the sample size n varies from 100 to 1000, for each $p \in \{50, 200\}$.

which specifies a sparse vector $\beta^* = (\beta_1^*, \dots, \beta_p^*)^\top \in \mathbb{R}^p$. Given each simulated dataset, the tuning parameter $\lambda > 0$ is chosen to minimize a 10-fold cross-validated prediction error.

Figure 1 summarizes the results of the treatment effect-modifier selection performance with respect to the true/false positive rates (the left/right two panels, respectively), comparing the proposed additive regression to the linear regression approach, which are reported as the averages (and ± 1 standard deviations) across the 200 simulation replications. Figure 1 illustrates that, for the both $p = 50$ and $p = 200$ cases, the proportion of the correctly selecting treatment effect-modifiers (i.e., the “true positive”; the left two panels) of the additive regression method (the red solid curves) tends to 1, with n increasing from $n = 100$ to $n = 1000$, while the proportion of incorrectly selecting treatment effect-modifiers (i.e., the “false positive”; the right two panels) tends to be bounded above by a small number. On the other hand, the proportion of correctly selecting treatment effect-modifiers for the linear regression method (the blue dotted curves) tends to be only around 0.5 for both choices of p . In Figure S.2 of Supplementary Materials available at *Biostatistics* online, we further examine the true positive rates reported in Figure 1, by separately displaying the true positive rates associated with selection of X_1, X_2 , and X_3 respectively. The more flexible additive regression significantly outperforms the linear regression in terms of selecting the covariates X_2 and X_3 , i.e., the ones that have the nonlinear interaction effects (see Figure S.1 of Supplementary Materials available at *Biostatistics* online for the functions associated with the interaction effects) with A , while the both methods perform at a similar level in selection of X_1 which has a linear interaction effect with A .

4.2. Individualized treatment rule estimation performance

In this subsection, we assess the optimal ITR estimation performance of the proposed method based on simulations. We generate a vector of covariates $\mathbf{X} = (X_1, \dots, X_p)^\top \in \mathbb{R}^p$, $p = 50$ (a $p = 100$ case is considered in Section S.5.1 of Supplementary Materials available at *Biostatistics* online) based on a multivariate normal distribution with each component having the marginal distribution $\mathcal{N}(0, (\pi/2)^2)$ with the correlation between the components $\text{corr}(X_j, X_k) = 0.1^{|j-k|}$. Responses were generated, for (i) “highly nonlinear” A -by- \mathbf{X} interactions:

$$Y = \delta \sum_{j=1}^5 \sin(X_j) + 2(A - 1.5) \{ \cos(X_1) - \cos(X_2) + \xi \sin(X_1 X_2) \} + \epsilon \quad A \in \{1, 2\}, \quad (4.22)$$

and for (ii) “*moderately nonlinear*” A -by- \mathbf{X} interactions:

$$Y = \delta \sum_{j=1}^5 \sin(X_j) + 2(A - 1.5) \left\{ \frac{2 \exp(5X_1)}{1 + \exp(5X_1)} - \frac{2 \exp(5X_2)}{1 + \exp(5X_2)} + \xi \frac{X_1 X_2}{2} \right\} + \epsilon \quad A \in \{1, 2\}, \quad (4.23)$$

where the treatment variable $A \in \{1, 2\}$ is generated independently from the covariates \mathbf{X} and the error term $\epsilon \sim \mathcal{N}(0, 0.5^2)$, such that $\Pr(A = 1) = \Pr(A = 2) = 1/2$. Models (4.22) and (4.23) are indexed by a pair $\{\delta, \xi\}$. First, the parameter $\delta \in \{1, 2\}$ controls the proportion of the variance of the response Y attributable to the \mathbf{X} “main” effect: $\delta = 1$ corresponds to a *moderate* \mathbf{X} main effect contribution; $\delta = 2$ corresponds to a *large* \mathbf{X} main effect contribution. Estimation of the interaction effect becomes more difficult with a larger δ . Second, the parameter $\xi \in \{0, 1\}$ determines whether the A -by- \mathbf{X} interaction effect term has an exact additive regression structure ($\xi = 0$) or whether it deviates from an additive structure ($\xi = 1$). In the case of $\xi = 0$, the proposed model (2.1) is correctly specified, whereas, for the case of $\xi = 1$, it is misspecified. For each scenario, we consider the following four approaches to estimating the optimal ITR \mathcal{D}^{opt} in (3.19).

1. The proposed additive regression approach (2.6), estimated via Algorithm 1. The dimension of the basis function Ψ_j in (3.10) is taken to be $d_j = 6$ ($j = 1, \dots, p$). Given estimates $\{\hat{g}_j^*\}$, the estimate of \mathcal{D}^{opt} in (3.19) is $\hat{\mathcal{D}}^{\text{opt}}(\mathbf{X}) = \arg \max_{a \in \{1, \dots, L\}} \sum_{j=1}^p \hat{g}_{j,a}^*(X_j)$.
2. The linear regression (MC) approach (4.21) of Tian and others (2014), implemented through the R-package `glmnet`, with the sparsity tuning parameter λ selected by minimizing a 10-fold cross-validated prediction error. Given an estimate $\hat{\beta}^* = (\hat{\beta}_1^*, \dots, \hat{\beta}_p^*)^\top$, the corresponding estimate of \mathcal{D}^{opt} in (3.19) is $\hat{\mathcal{D}}^{\text{opt}}(\mathbf{X}) = \arg \max_{a \in \{1, \dots, L\}} \sum_{j=1}^p (a - 1.5) X_j \hat{\beta}_j^*$.
3. The OWL method (Zhao and others, 2012) based on a Gaussian radial kernel, implemented in the R-package `DTRlearn`, with a set of feature transformed (FT) covariates $\{\hat{g}_{j,1}^*(X_j), j = 1, \dots, p\}$ used as an input to the OWL method, in which the functions $\hat{g}_{j,1}^*(\cdot)$ ($j = 1, \dots, p$) are obtained from the approach in 1. To improve the efficiency of the OWL, we employ the augmented OWL approach of Liu and others (2018). The inverse bandwidth parameter σ_n^2 and the tuning parameter κ in Zhao and others (2012) are chosen from the grid of $(0.01, 0.02, 0.04, \dots, 0.64, 1.28)$ and that of $(0.25, 0.5, 1, 2, 4)$ (the default setting of `DTRlearn`), respectively, based on a 5-fold cross-validation.
4. The same (OWL) approach as in 3 but based on the original features $\{X_j, j = 1, \dots, p\}$.

For each simulation run, we estimate \mathcal{D}^{opt} from each of the four methods based on a training set (of size $n \in \{250, 500\}$), and for evaluation of these methods, we compute the value $V(\hat{\mathcal{D}}^{\text{opt}})$ in (3.18) for each estimate $\hat{\mathcal{D}}^{\text{opt}}$, using a Monte Carlo approximation based on a random sample of size 10^3 . Since we know the true data generating model in simulation studies, the optimal \mathcal{D}^{opt} can be determined for each simulation run. Given each estimate $\hat{\mathcal{D}}^{\text{opt}}$ of \mathcal{D}^{opt} , we report $V(\hat{\mathcal{D}}^{\text{opt}}) - V(\mathcal{D}^{\text{opt}})$, as the performance measure of $\hat{\mathcal{D}}^{\text{opt}}$. A larger value (i.e., a smaller difference from the optimal value) of the measure indicates better performance.

In Figure 2, we present the boxplots, obtained from 100 simulation runs, of the normalized values $V(\hat{\mathcal{D}}^{\text{opt}})$ (normalized by the optimal values $V(\mathcal{D}^{\text{opt}})$) of the decision rules $\hat{\mathcal{D}}^{\text{opt}}$ estimated from the four approaches, for each combination of $n \in \{250, 500\}$, $\xi \in \{0, 1\}$ (corresponding to *correctly specified* or *misspecified* additive interaction effect models, respectively) and $\delta \in \{1, 2\}$ (corresponding to *moderate* or *large* main effects, respectively), for the *highly nonlinear* and *moderately nonlinear* A -by- \mathbf{X} interaction effect scenarios, in the top and bottom panels, respectively. The proposed additive regression clearly outperforms the OWL (without feature transformation) method in all scenarios (both the top and bottom

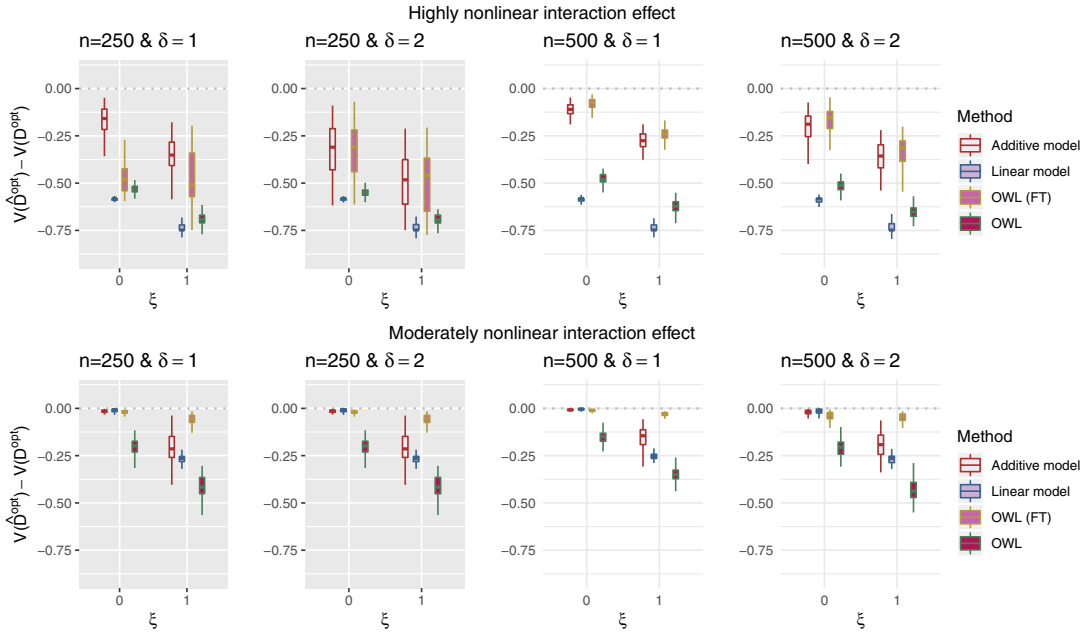


Fig. 2. Boxplots based on 100 simulation runs, comparing the four approaches to estimating \mathcal{D}^{opt} , given each scenario indexed by $\xi \in \{0, 1\}$, $\delta \in \{1, 2\}$, and $n \in \{250, 500\}$, for the *highly nonlinear* A -by- X interaction effect case in the top panels, and the *moderately nonlinear* A -by- X interaction effect case in the bottom panels. The dotted horizontal line represents the optimal value corresponding to \mathcal{D}^{opt} .

panels) and also the linear regression approach in all of the *highly nonlinear* A -by- X interaction effect scenarios (the top panels). For the *moderately nonlinear* A -by- X interaction effect scenarios (the bottom panels), when $\xi = 0$, all the methods except the OWL perform at a near-optimal level. On the other hand, when $\xi = 1$ (i.e., when the underlying A -by- X interaction effect model deviates from the additive structure), the more flexible additive model significantly outperforms the linear model. We have also considered a *linear* A -by- X interaction effect case in Section S.5.2 of Supplementary Materials available at *Biostatistics* online, in which the linear regression outperforms the additive regression, but only slightly, whereas if the underlying model deviates from the exact linear structure and $n = 500$, the more flexible additive regression tends to outperform the linear model. This suggests that, in the absence of prior knowledge about the form of the interaction effect, employing the proposed additive regression is more suitable for optimizing ITRs than the linear regression. Comparing the two OWL methods (OWL (FT) and OWL), Figure 2 illustrates that the feature transformation based on the estimated component functions $\{\hat{g}_j^*, j = 1, \dots, p\}$ provides a considerable benefit in their performance. This suggests the utility of the proposed model (2.1) as a potential feature transformation and selection tool for a machine learning algorithm for optimizing ITRs. In Figure 2, comparing the cases with $\delta = 2$ to those with $\delta = 1$, the increased magnitude of the main effect generally dampens the performance of all approaches, as the “noise” variability in the data generation model increases.

5. APPLICATION TO DATA FROM A DEPRESSION CLINICAL TRIAL

In this section, we illustrate the utility of the proposed additive regression for estimating treatment effect modification and optimizing individualized treatment rules, using data from a depression clinical trial

study, comparing an antidepressant and placebo for treating major depressive disorder (Trivedi and others, 2016). The goal of the study is to identify baseline characteristics that are associated with differential response to the antidepressant versus placebo and to use those characteristics to guide treatment decisions when a patient presents for treatment.

Study participants (a total of $n = 166$ participants) were randomized to either placebo ($A = 1; n_1 = 88$) or an antidepressant (sertraline) ($A = 2; n_2 = 78$). Subjects were monitored for 8 weeks after initiation of treatment, and the primary endpoint of interest was the Hamilton Rating Scale for Depression (HRSD) score at week 8. The outcome Y was taken to be the improvement in symptoms severity from baseline to week 8, taken as the difference, i.e., we take: week 0 HRSD score–week 8 HRSD score. (Larger values of the outcome Y are considered desirable.) The study collected baseline patient clinical data, prior to treatment assignment. These pretreatment clinical data $\mathbf{X} = (X_1, X_2, \dots, X_{13})^\top \in \mathbb{R}^{13}$ include: $X_1 =$ Age at evaluation; $X_2 =$ Severity of depressive symptoms measured by the HRSD at baseline; $X_3 =$ Logarithm of duration (in month) of the current major depressive episode; and $X_4 =$ Age of onset of the first major depressive episode. In addition to these standard clinical assessments, patients underwent neuropsychiatric testing at baseline to assess psychomotor slowing, working memory, reaction time (RT), and cognitive control (e.g., post-error recovery), as these behavioral characteristics are believed to correspond to biological phenotypes related to response to antidepressants (Petkova and others, 2017) and are considered as potential modifiers of the treatment effect. These neuropsychiatric baseline test measures include: $X_5 =$ (A not B) RT-negative; $X_6 =$ (A not B) RT-non-negative; $X_7 =$ (A not B) RT-all; $X_8 =$ (A not B) RT-total correct; $X_9 =$ Median choice RT; $X_{10} =$ Word fluency; $X_{11} =$ Flanker accuracy; $X_{12} =$ Flanker RT; $X_{13} =$ Post-conflict adjustment.

The proposed approach (2.6) to estimating the A -by- \mathbf{X} interaction effect part $\sum_{j=1}^p g_{j,A}^*(X_j)$ of model (2.1), estimated via Algorithm 1, simultaneously selected 3 pretreatment covariates as treatment effect-modifiers: X_1 (“Age at evaluation”), X_{10} (“Word fluency test”), and X_{11} (“Flanker accuracy test”). The top panels in Figure 3 illustrate the estimated non-zero component functions $\{\hat{g}_j^* \neq 0, j = 1, \dots, 13\}$ (i.e., the component functions corresponding to the selected covariates X_1, X_{10} , and X_{11}) and the associated partial residuals. The linear regression approach (4.21) to estimating the A -by- \mathbf{X} interactions selected a single covariate, X_{11} , as a treatment effect-modifier.

Note, in the binary (i.e., $L = 2$) treatment case, any ITR $\mathcal{D}(\mathbf{X})$ partitions the domain of \mathbf{X} , \mathbb{R}^p , into two regions: $R_1 = \{\mathbf{x} \in \mathbb{R}^p : \mathcal{D}(\mathbf{x}) = 1\}$ and $R_2 = \{\mathbf{x} \in \mathbb{R}^p : \mathcal{D}(\mathbf{x}) = 2\}$. Let $\hat{\beta}_{0,a}$ ($a = 1, 2$) represent the treatment a -specific intercept estimates. Then the proposed estimator $\hat{\mathcal{D}}^{\text{opt}}(\mathbf{X}) = \arg \max_{a \in \{1,2\}} \{\hat{\beta}_{0,a} + \sum_{j=1}^p \hat{g}_{j,a}^*(X_j)\}$, is equivalent to:

$$\hat{\mathcal{D}}^{\text{opt}}(\mathbf{X}) = I \left[\hat{\beta}_{0,2} - \hat{\beta}_{0,1} + \sum_{j=1}^p \{\hat{g}_{j,2}^*(X_j) - \hat{g}_{j,1}^*(X_j)\} > 0 \right] + 1. \quad (5.24)$$

In this dataset, $\hat{\beta}_{0,1} = 6.34$ (corresponding to the placebo) and $\hat{\beta}_{0,2} = 7.5$ (corresponding to the drug). The optimal treatment region in the X_1, X_{10}, X_{11} space, implied by the ITR (5.24), is illustrated in Figure 4, where we have utilized $40 \times 40 \times 40$ equally spaced grid points (i.e., 40 for each axis) for visualization of the regions R_1 and R_2 .

For an alternative way of visualizing the ITR (5.24), let us define a 1D index $h(\mathbf{X}) = \sum_{j=1}^p \hat{g}_{j,1}^*(X_j)$. By the constraint (2.2), we have the relationship $\sum_{j=1}^p \hat{g}_{j,2}^*(X_j) = -\hat{\pi}_2^{-1} \hat{\pi}_1 \{\sum_{j=1}^p \hat{g}_{j,1}^*(X_j)\}$. Therefore, the term $\sum_{j=1}^p \{\hat{g}_{j,2}^*(X_j) - \hat{g}_{j,1}^*(X_j)\}$ in the decision rule (5.24) can be reparametrized, with respect to $h(\mathbf{X})$, as $(-\hat{\pi}_2^{-1} \hat{\pi}_1 - 1)h(\mathbf{X})$. In this dataset, $\hat{\pi}_1 = 0.53$ and $\hat{\pi}_2 = 0.47$, and thus the ITR (5.24) can be re-written as: $\hat{\mathcal{D}}^{\text{opt}}(\mathbf{X}) = I[1.16 - 2.12h(\mathbf{X}) > 0] + 1$; this ITR indicates that, for a patient with pretreatment characteristics \mathbf{x} , if $h(\mathbf{x}) < 0.55$ ($\approx 1.16/2.12$), then he/she would be recommended the active drug

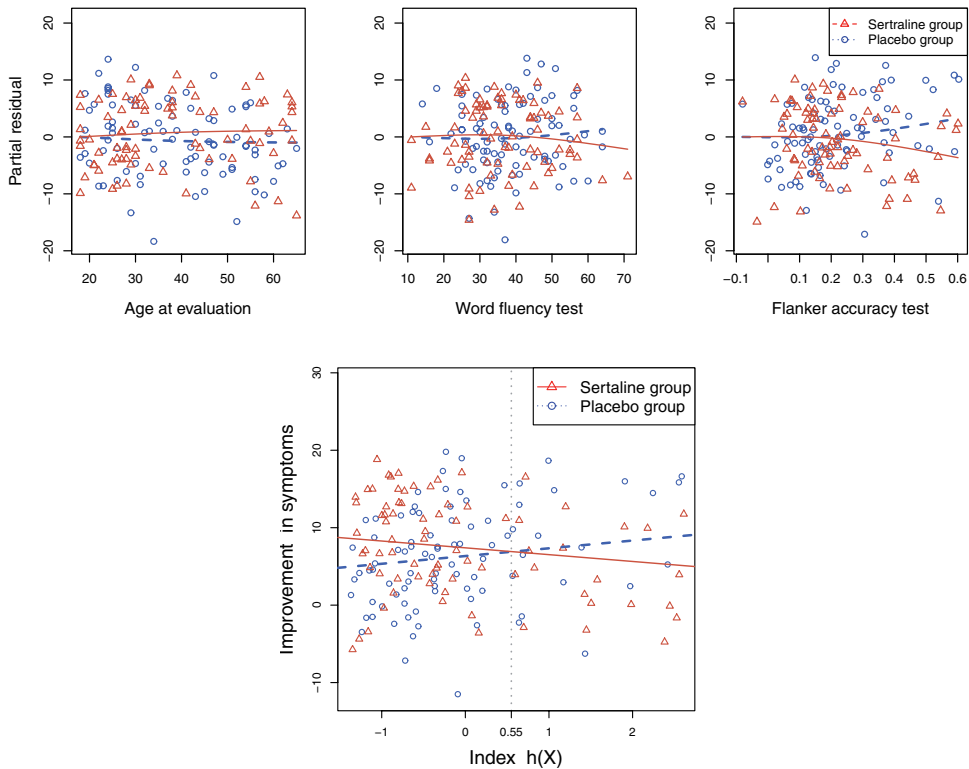


Fig. 3. The top panels: Scatterplots of partial residuals vs. the covariates associated with estimated non-zero component functions $\{\hat{g}_j^* \neq 0\}$ for placebo (blue circles) and the active drug (red triangles) treated participants; for each panel, the blue dashed curve represents $\hat{g}_{j,1}^*(\cdot)$, corresponding to the placebo ($a = 1$), and the red solid curve represents $\hat{g}_{j,2}^*(\cdot)$, corresponding to the active drug ($a = 2$). The bottom panel: A scatterplot of the treatment response y versus the $index\ h(x) = \sum_{j=1}^p \hat{g}_{j,1}^*(x_j)$, for the drug (red triangles) and the placebo (blue circles) groups; the blue dashed line is $y = \hat{\beta}_{0,1} + h(x)$ and the red solid line is $y = \hat{\beta}_{0,2} - \hat{\pi}_2^{-1} \hat{\pi}_1 h(x)$; a gray dashed vertical line indicates the threshold value $h(x) = 0.55$ associated with the ITR $\hat{D}^{opt}(x)$.

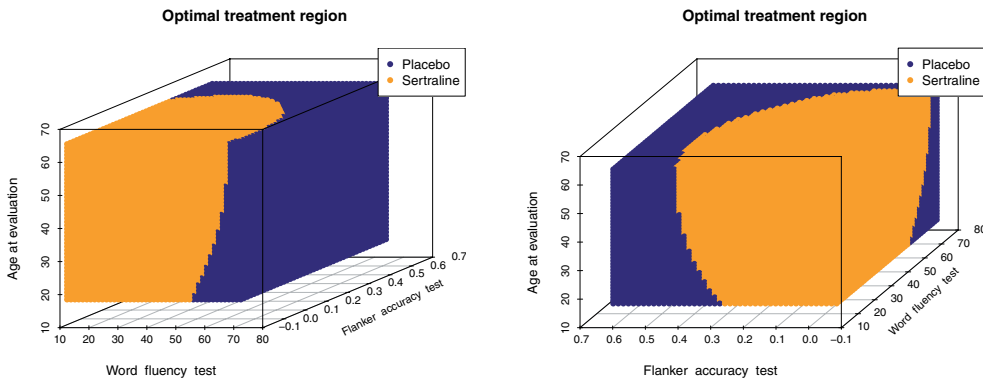


Fig. 4. The decision regions R_1 (corresponding to the placebo, dark blue) and R_2 (corresponding to the active drug, orange) displayed in the 3D cube of the three covariates X_1 (age at evaluation), X_{10} (word fluency test), and X_{11} (Flanker accuracy test), evaluated at $40 \times 40 \times 40$ equally spaced grid points (40 for each axis), shown from two different angles (left and right panels).




Fig. 5. Boxplots of the estimated values of the treatment rules \hat{D}^{opt} estimated from 6 approaches, obtained from 500 randomly split testing sets. Higher values are preferred.

($a = 2$) and the placebo ($a = 1$), otherwise. For example, for a patient with $x_1 = 40$, $x_{10} = 50$, and $x_{11} = 0.4$, the *index* $h(\mathbf{x}) = 1.01 > 0.55$ (see the bottom panel in Figure 3 for a visualization) and thus the patient would be recommended the placebo.

To evaluate the performance of the ITRs (\hat{D}^{opt}) obtained from the four different approaches described in Section 4.2, we randomly split the data into a training set and a testing set (of size \tilde{n}) using a ratio of 5 to 1, replicated 500 times, each time computing an ITR \hat{D}^{opt} based on the training set, then estimating its value $V(\hat{D}^{\text{opt}})$ in (3.18) by an inverse probability weighted estimator (Murphy, 2005): $\hat{V}(\hat{D}^{\text{opt}}) = \sum_{i=1}^{\tilde{n}} Y_i I_{(A_i = \hat{D}^{\text{opt}}(X_i))} / \sum_{i=1}^{\tilde{n}} I_{(A_i = \hat{D}^{\text{opt}}(X_i))}$, computed based on the testing set of size \tilde{n} . For comparison, we also include two naïve rules: treating all patients with placebo (“All PBO”) and treating all patients with the active drug (“All DRUG”), each regardless of the individual patient’s characteristics \mathbf{X} . The resulting boxplots obtained from the 500 random splits are illustrated in Figure 5. A larger value of the measure indicates better performance.

The results in Figure 5 demonstrate that the proposed additive regression approach, which allows non-linear flexibility in developing ITRs, tends to outperform the linear regression approach, in terms of the estimated value. The additive regression approach also shows some superiority over the method OWL (without feature transformation). In comparison to the OWL methods, the proposed additive regression, in addition to its superior computational efficiency, provides a means of simultaneously selecting treatment effect-modifiers and allows a visualization for the heterogeneous effects attributable to each estimated treatment effect-modifier as in Figure 3, which is an appealing feature in practice. Moreover, the estimated component functions $\{\hat{g}_j^*, j = 1, \dots, p\}$ of the proposed regression provide an effective means of performing feature transformation for $\{X_j, j = 1, \dots, p\}$. As in Section 4.2, the FT OWL approach appears to have a considerable improvement over the OWL that bases on the original untransformed covariates.

6. DISCUSSION

In this article, we have developed a sparse additive model, via a structural constraint, specifically geared to identify and model treatment effect-modifiers. The approach utilizes an efficient back-fitting algorithm for model estimation and variable selection. The proposed sparse additive model for treatment effect

modification extends existing linear model-based regression methods by providing nonlinear flexibility to modeling treatment-by-covariate interactions. Encouraged by our simulation results and the application, future work will investigate the asymptotic properties related to treatment effect-modifier selection and estimation consistency, in addition to developing hypothesis testing procedures for treatment-by-covariates interaction effects.

Modern advances in biotechnology, using measures of brain structure and function obtained from neuroimaging modalities (e.g., magnetic resonance imaging (MRI), functional MRI, and electroencephalography), show the promise of discovering potential biomarkers for heterogeneous treatment effects. These high-dimensional data modalities are often in the form of curves or images and can be viewed as functional data (e.g., Ramsay and Silverman, 1997). Future work will also extend the additive model approach to the context of functional additive regression (e.g., Fan and others, 2014, 2015). The goal of these extensions will be to handle a large number of functional-valued covariates while achieving simultaneous variable selection, which will extend current functional linear model-based methods for precision medicine (McKeague and Qian, 2014; Ciarleglio and others, 2015, 2018) to a more flexible functional regression setting, as well as to longitudinally observed functional data (e.g., Park and Lee 2019).

7. SOFTWARE

R-package `samTEMsel` (Sparse Additive Models for Treatment Effect-Modifier Selection) contains R-codes to perform the methods proposed in the article, and is publicly available on GitHub (github.com/syhyunpark/samTEMsel).

SUPPLEMENTARY MATERIALS

Supplementary material is available at <http://biostatistics.oxfordjournals.org>

FUNDING

National Institute of Mental Health (NIMH) (5 R01 MH099003).

ACKNOWLEDGMENTS

We are grateful to the editors, the associate editor, and two referees for their insightful comments and suggestions.

Conflict of Interest: None declared.

REFERENCES

- ASHLEY, E. (2015). The precision medicine initiative: a new national effort. *The Journal of the American Medical Association* **313**, 2117.
- CIARLEGLIO, A., PETKOVA, E., OGDEN, R. T. AND TARPEY, T. (2015). Treatment decisions based on scalar and functional baseline covariate decisions based on scalar and functional baseline covariates. *Biometrics* **71**(4), 884–894.
- CIARLEGLIO, A., PETKOVA, E., OGDEN, R. T. AND TARPEY, T. (2018). Constructing treatment decision rules based on scalar and functional predictors when moderators of treatment effect are unknown. *Journal of Royal Statistical Society: Series C* **67**, 1331–1356.
- FAN, Y., FOUTZ, N., JAMES, G. M. AND JANK, W. (2014). Functional response additive model estimation with online virtual stock markets. *The Annals of Applied Statistics* **8**, 2435–2460.
- FAN, Y., JAMES, G. M. AND RADCHANKO, P. (2015). Functional additive regression. *The Annals of Statistics* **43**, 2296–2325.

- FERNANDES, B., WILLIAMS, L., STEINER, J., LEBOYER, M., CARVALHO, A. AND BERK, M. (2017). The new field of 'precision psychiatry'. *BMC Medicine* **15**, 80.
- HASTIE, T. AND TIBSHIRANI, R. (1999). *Generalized Additive Models*. London: Chapman and Hall/CRC.
- JENG, X., LU, W. AND PENG, H. (2018). High-dimensional inference for personalized treatment decision. *Electronic Journal of Statistics* **12**, 2074–2089.
- KANG, C., JANES, H. AND HUANG, Y. (2014). Combining biomarkers to optimize patient treatment recommendations. *Biometrics* **70**, 696–707.
- KOSOROK, M. R. AND LABER, E. B. (2019). Precision medicine. *Annual Review of Statistics and Its Application* **6**(1), 263–286.
- KUHN, M. AND JOHNSON, K. (2019). *Feature Engineering and Selection: A Practical Approach for Predictive Models*. London: Chapman and Hall/CRC.
- LABER, E. B. AND ZHAO, Y. (2015). Tree-based methods for individualized treatment regimes. *Biometrika* **102**, 501–514.
- LIU, Y., WANG, Y., KOSOROK, M. R., ZHAO, Y. AND ZENG, D. (2018). Augmented outcome-weighted learning for estimating optimal dynamic treatment regimens. *Statistics in Medicine* **37**, 3776–3788.
- LU, W., ZHANG, H. AND ZENG, D. (2011). Variable selection for optimal treatment decision. *Statistical Methods in Medical Research* **22**, 493–504.
- MCKEAGUE, I. AND QIAN, M. (2014). Estimation of treatment policies based on functional predictors. *Statistica Sinica* **24**, 1461–1485.
- MURATA, N. AND AMARI, S. (1994). Network Information Criterion-determining the number of hidden units for an artificial neural network model. *IEEE Transactions on Neural Networks* **5**(6), 865–872.
- MURPHY, S. A. (2003). Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **65**, 331–355.
- MURPHY, S. A. (2005). A generalization error for q-learning. *Journal of Machine Learning* **6**, 1073–1097.
- Park, H. and Lee, S. (2019). Logistic regression error-in-covariate models for longitudinal high-dimensional covariates. *Stat* **8**, e246.
- Park, H., Petkova, E., Tarpey, T. and Ogden, R. T. (2020). A constrained single-index regression for estimating interactions between a treatment and covariates. *Biometrics*, 1–13. <https://doi.org/10.1111/biom.13320>.
- PETKOVA, E., OGDEN, R. T., TARPEY, T., CIARLEGLIO, A., JIANG, B., SU, Z., CARMODY, T., ADAMS, P., KRAEMER, H., GRANNEMANN, B. and others. (2017). Statistical analysis plan for stage 1 embark (establishing moderators and biosignatures of antidepressant response for clinical care) study. *Contemporary Clinical Trials Communications* **6**, 22–30.
- QIAN, M. AND MURPHY, S. A. (2011). Performance guarantees for individualized treatment rules. *The Annals of Statistics* **39**(2), 1180–1210.
- RAMSAY, J. O. AND SILVERMAN, B. W. (1997). *Functional Data Analysis*. New York: Springer.
- RAVIKUMAR, P., LAFFERTY, J., LIU, H. AND WASSERMAN, L. (2009). Sparse additive models. *Journal of Royal Statistical Society: Series B* **71**, 1009–1030.
- ROYSTON, P. AND SAUERBREI, W. (2008). Interactions between treatment and continuous covariates: a step toward individualizing therapy. *Journal of Clinical Oncology* **26**(9), 1397–1399.
- SHI, C., SONG, R. AND LU, W. (2016). Robust learning for optimal treatment decision with np-dimensionality. *Electronic Journal of Statistics* **10**, 2894–2921.

- SONG, R., KOSOROK, M., ZENG, D., ZHAO, Y., LABER, E. B. AND YUAN, M. (2015). On sparse representation for optimal individualized treatment selection with penalized outcome weighted learning. *Stat* **4**, 59–68.
- TIAN, L., ALIZADEH, A., GENTLES, A. AND TIBSHRANI, R. (2014). A simple method for estimating interactions between a treatment and a large number of covariates. *Journal of the American Statistical Association* **109**(508), 1517–1532.
- TRIVEDI, M., MCGRATH, P., FAVA, M., PARSEY, R., KURIAN, B., PHILLIPS, M., PQUENDO, M., BRUDER, G., PIZZAGALLI, D., TOUPS, M., *and others*. (2016). Establishing moderators and biosignatures of antidepressant response in clinical care (embarc): rationale and design. *Journal of Psychiatric Research* **78**, 11–23.
- TSENG, P. (2001). Convergence of block coordinate descent method for nondifferentiable maximization. *Journal of Optimization Theory and Applications* **109**, 474–494.
- ZHANG, B., TSIATIS, A. A., DAVIDIAN, M., ZHANG, M. AND LABER, E. (2012). Estimating optimal treatment regimes from classification perspective. *Stat* **1**, 103–114.
- ZHAO, Y., LABER, E., NING, Y., SAHA, S. AND SANDS, B. (2019). Efficient augmentation and relaxation learning for individualized treatment rules using observational data. *Journal of Machine Learning Research* **20**, 1–23.
- ZHAO, Y., ZENG, D., RUSH, A. J. AND KOSOROK, M. R. (2012). Estimating individualized treatment rules using outcome weighted learning. *Journal of the American Statistical Association* **107**, 1106–1118.
- ZHAO, Y., ZHENG, D., LABER, E. B. AND KOSOROK, M. R. (2015). New statistical learning methods for estimating optimal dynamic treatment regimes. *Journal of the American Statistical Association* **110**, 583–598.

[Received March 17, 2020; revised July 5, 2020; accepted for publication July 10, 2020]