Published in final edited form as:

Biometrics. 2015 September ; 71(3): 645–653. doi:10.1111/biom.12322.

Regularized Outcome Weighted Subgroup Identification for Differential Treatment Effects

Yaoyao Xu¹, Menggang Yu^{2,*}, Ying-Qi Zhao², Qiefeng Li³, Sijian Wang^{1,2}, and Jun Shao¹

¹Department of Statistics, University of Wisconsin, Madison, Wisconsin, U.S.A

²Department of Biostatistics & Medical Informatics, University of Wisconsin, Madison, Wisconsin, U.S.A

³Department of Operation Research and Financial Engineering, Princeton University, Princeton, New Jersey, U.S.A

Summary

To facilitate comparative treatment selection when there is substantial heterogeneity of treatment effectiveness, it is important to identify subgroups that exhibit differential treatment effects. Existing approaches model outcomes directly and then define subgroups according to interactions between treatment and covariates. Because outcomes are affected by both the covariate–treatment interactions and covariate main effects, direct modeling outcomes can be hard due to model misspecification, especially in presence of many covariates. Alternatively one can directly work with differential treatment effect estimation. We propose such a method that approximates a target function whose value directly reflects correct treatment assignment for patients. The function uses patient outcomes as weights rather than modeling targets. Consequently, our method can deal with binary, continuous, time-to-event, and possibly contaminated outcomes in the same fashion. We first focus on identifying only directional estimates from linear rules that characterize important subgroups. We further consider estimation of comparative treatment effects for identified subgroups. We demonstrate the advantages of our method in simulation studies and in analyses of two real data sets.

Keywords

Comparative effectiveness; Heterogeneity of treatment effectiveness; Regularization; Subgroup; Variable selection

1. Introduction

Subgroup identification for differential outcomes is becoming an increasingly important methodological research topic in this era of personalized medicine and high medical costs (Ruberg, Chen, and Wang, 2010). Identification of subgroups via valid predictive markers

* meyu@biostat.wisc.edu.

Supplementary Materials

Web Appendices and program codes, referenced in Sections 2–5, are available with this paper at the *Biometrics* website on Wiley Online Library.

can facilitate tailored intervention to maximize effectiveness. The controversy and challenges of subgroup identification have been well documented (Brookes et al., 2004; Lagakos, 2006; Ruberg et al., 2010). Yet there are still great needs for descriptive and exploratory subgroup identification using existing data to develop new hypothesis, to facilitate study population selection for further investigation, and to generate new evidence for medical decision making (Varadhan et al., 2013).

The classical approach to identify subgroups involves fitting of a parametric outcome model which includes covariate main effects as well as interactions between covariates and treatment. Subgroups are identified based on significance of specific interaction terms (Kehl and Ulm, 2006). Besides the multiplicity issue, misspecification of the main effects in the presence of many covariates interferes with the covariate–treatment interaction estimation and thus impedes valid subgroup identification. Nonparametric approaches based on the classification and regression tree methodology (Breiman et al., 1984) were attempted to separate the main effects from the covariate–treatment interaction effects, either through multiple testing strategies (Su et al., 2008, 2009; Lipkovich et al., 2011; He, 2012) or prediction (Foster, Taylor, and Ruberg, 2011). For example, Su et al. (2009) developed an interaction tree method by building splitting rules based on covariate–treatment interaction tests. Foster et al. (2011) developed the virtual twins method by first estimating treatment differences based on random forests and then applying tree to those estimates to form subgroups. Alternatively, Cai et al. (2011) built in calibration steps to alleviate the model misspecification problem when using parametric or semiparametric models.

In this article, we propose an approach that works directly with differential treatment effect estimation. Instead of modeling the outcome, we use a causal inference formulation (Zhao et al., 2012) that leads to a target function directly reflecting the covariate–treatment interaction effect. This spares the need to model the covariate main effects. The function uses patient outcomes as weights rather than modeling targets. Consequently, our method can deal with binary, continuous, time-to-event, and possibly contaminated outcomes in the same fashion, and can be implemented using existing packages. Foster et al. (2011) pointed out that instead of focusing on finding the optimal or largest subgroup with enhanced treatment effects, it may be more desirable to determine a sub-optimal subgroup from simple and interpretable predictive variables. Similar to Imai and Ratkovic (in press), we focus on linear rules and build in smoothness constraints that allow the resulting subgroup to be more interpretable. Recent applications often involve a large number of covariates. Considering covariate–treatment interaction terms also increases the dimension of the covariate matrix. Thus, we add a regularization step in the subgroup identification process. The framework and details of our methodology are given in Section 2. We note that such regularization has been used in dynamic treatment regimens (Song et al., in press).

Our motivating example is a mammography screening randomized trial that tested the efficacy of interventions to promote mammography screening among eligible women who were 51 years or older (Champion et al., 2007). Data were collected from 09/01/96 to 11/30/02. Different from the usual drug trial, this study intended to change health behavior of the study subjects, which was more complicated and challenging. Besides establishing an overall efficacy, it was also highly desirable to determine whether the interventions were

more effective for certain subgroups determined by their social and behavioral variables. For this reason, in addition to demographics and social economic variables, the study collected many health behavioral variables such as self-efficacy, susceptibility, knowledge, fear, fatalism, benefits, and barriers based on the well-established Health Belief Model (Janz and Becker, 1984). These variables intended to capture corresponding behavioral constructs. If interventions benefit subjects differently according to their characteristics, then they can be applied more cost-effectively. In addition, newer interventions can be designed to target relevant behavioral constructs and other modifiable characteristics. It is therefore highly desirable that the resulting treatment assignment rules are interpretable and simple. For this purpose, we focus on finding one-way and two-way interactions among the covariates and interventions. Nevertheless, a large number of variables still need to be investigated due to the many levels of the behavioral variables, prompting us to focus on parsimonious and interpretable models.

In Section 2, we argue that for many popular models, it suffices to find the optimal *linear* rule. The linearity allows regularization that can lead to sparse solutions, in contrast with Zhao et al. (2012). In Section 3, we develop an inference procedure on the enhanced treatment effects to further evaluate the benefit of the identified subgroup. Such a step is important to aid practical decision making. Our method is compared with some other recent subgroup identification methods, including a sparse version of Zhao et al. (2012), through simulation studies in Section 4, and is applied to two real life examples in Section 5. We conclude with a discussion of our findings and future work in Section 6.

2. Subgroup Identification

Assume data are collected from a randomized trial with a binary treatment $T \in \{-1, 1\}$ being according to $P(T=1) = \pi$ a clinical outcome Y , and a p -dimensional covariates vector $\mathbf{X} = (X_1, \dots, X_p)'$. We allow $\mathbf{X}$ to contain derived covariates such as interaction terms as well as dummy variables for categorical covariates. Let $\mathcal{D}(\mathbf{X})$ be an assignment rule based on $\mathbf{X}$, P the distribution of $(Y, T, \mathbf{X})$ in the data and $P^{\mathcal{D}}$ the distribution of $(Y, T, \mathbf{X})$ given that $T = \mathcal{D}(\mathbf{X})$. Then, the expected outcome $E^{\mathcal{D}}(Y)$ under the rule $\mathcal{D}$ is given by (Qian and Murphy, 2011; Zhao et al., 2012; Zhang et al., 2012)

$$E^{\mathcal{D}}(Y) = \int Y dP^{\mathcal{D}} = \int Y \frac{dP^{\mathcal{D}}}{dP} dP = E \left[\frac{I(T = \mathcal{D}(\mathbf{X}))}{T\pi + (1-T)/2} Y \right],$$

where $I(\cdot)$ is the indicator function. Our main goal is to assign each individual to an appropriate treatment based on $\mathbf{X}$ to optimize the clinical outcome. Assuming larger Y is preferable, we want to find the optimal rule $\mathcal{D}^*$ that maximizes $E^{\mathcal{D}}(Y)$, or equivalently

$$\mathcal{D}^*(\mathbf{X}) = \arg \min_{\mathcal{D}} E \left[\frac{I(T \neq \mathcal{D}(\mathbf{X}))}{T\pi + (1-T)/2} Y \right]. \quad (1)$$

In general, the optimal rule $\mathcal{D}^*(X)$ could be a very complicated function of X . However, we show in the following proposition that, under mild conditions, the optimal rule $\mathcal{D}^*$ only depends on the sign of a *linear* function of X .

Proposition 1

Suppose $E(Y|T, X) = h(X, TX' \beta^*)$, where for any real number $a \in R$, the bivariate function $h(a, b)$ satisfies $h(a, b) = h(a, -b)$ for $b \geq 0$ and $h(a, b) < h(a, -b)$ for $b < 0$. Then, for almost surely any fixed x , $\mathcal{D}^*(x) = \text{sign}(x' \beta^*)$, where $\text{sign}(x) = 1$ if $x \geq 0$ and $\text{sign}(x) = -1$ otherwise.

Indeed, many popular statistical models satisfy the requirements of Proposition 1, including the following two classes of models.

Example 1

The response Y follows a single index model, $E(Y|T, X) = h\{\ell(X) + g(TX' \beta^*)\}$, where $h(x)$ and $g(x)$ are increasing functions. This example satisfies the assumptions in Proposition 1 and includes the logistic model as a special case. The optimal rule $\mathcal{D}^*$ is unrelated to $\ell(X)$, the main effects of X , which may not be linear in X .

Example 2

The response Y is a time-to-event variable and follows the Aalen-Cox model $\lambda_Y(y|X, T) = \lambda_0(y, X) \exp(-TX' \beta^*)$. In right-censored data, Y is censored by another variable C where C is independent of (Y, T) given X . That is we only observe $Y \wedge C = \min(Y, C)$. Then

$E(Y \wedge C|T, X) = \int e^{-\Lambda_0(u, X) \exp(-TX' \beta^*)} G(u|X) du$, where $G(u|X)$ is the conditional survival distribution of C given X . Therefore, by considering $Y \wedge C$ as the outcome, we have $h(a, b) = \int e^{-\Lambda_0(u, a) \exp(-b)} G(u|X) du$ and h satisfies the conditions in Proposition 1.

Obviously, one way to find the optimal rule is to estimate β^* by modeling $E(Y|T, X) = h(X, TX' \beta^*)$ directly. However, unless the form of h is simple and the dimension of X is low, this outcome modeling based approach may not be preferable as model misspecification can be severe. On the other hand, as implied by Proposition 1, correct treatment selection may not depend so much on the outcome model. Important variables for predicting the outcome also may not be relevant for treatment selection (Janes et al., 2011; Sargent and Mandrekhar, 2013). Even if h is a complicated function, the resulting treatment assignment rule can still be very simple and interpretable. For this reason, in the rest of this article we focus on finding a linear rule $\mathcal{D}^*(x) = \text{sign}(x' \beta^*)$ with

$$\beta^* \in \arg \min_{\beta \in \mathcal{R}^p} E \left[\frac{I\{T \neq \text{sign}(X' \beta)\}}{\pi T + (1 - T)/2} Y \right]. \quad (2)$$

Under the assumptions of Proposition 1, β^* is one solution. If $k\beta^*$ is a solution, so is $k\beta^*$ for any positive constant k , since $\mathcal{D}^*(x) = \text{sign}(x' \beta^*) = \text{sign}(x' k\beta^*)$ for any $k > 0$.

The 0–1 loss in (2) is not only non-convex but also non-continuous. In practice, it is computationally hard to solve (2). Similar to the common practice in classification algorithms (Bartlett, Jordan, and McAuliffe, 2006), we propose to replace the 0–1 loss with a convex and continuous surrogate loss $\phi(\cdot)$ and solve the following optimization problem:

$$\tilde{\beta} = \arg \min_{\beta \in \mathcal{R}^p} E \left[\frac{\phi(TX'\beta)}{\pi T + (1-T)/2} Y \right]. \quad (3)$$

Since (3) is a convex optimization problem, $\tilde{\beta}$ is its unique solution. We approximate the optimal rule $\mathcal{D}^*(x) = \text{sign}(x'\beta^*)$. In general, $\tilde{\beta}$ can differ from β^* due to the surrogate loss $\phi(\cdot)$. In some cases, $\tilde{\beta}$ and β^* differ only by a multiplier and, hence, $\text{sign}(X'\tilde{\beta}) = \text{sign}(X'\beta^*)$. One example is when $E(Y|T, X) = h(TX'\beta^*)$ for some function h and when the conditional expectation $E(Xb|X'\beta^*)$ is linear in $X'\beta^*$ for any $b \in \mathcal{R}^p$. This latter condition is satisfied when the distribution of X is elliptically symmetric, for example the normal distribution (Li, 1991). Even if the distribution of X is not exactly elliptically symmetric, Li (1991) argued that for most directions b , $E(Xb|X'\beta^*)$ is approximately linear in $X'\beta^*$. Hence $\tilde{\beta}$ is expected to be approximately parallel to β^* , which makes $\text{sign}(X'\tilde{\beta})$ a good approximation of $\text{sign}(X'\beta^*)$.

In applications, we solve (3) with expectation being replaced by the sample average based on independent samples (Y_i, T_i, X_i) , $i = 1, \dots, n$, identically distributed as (Y, T, X) . The dimension of β in (3) could be very high, especially when we include interactions between covariates. Hence, we propose to add two sparsity induced penalties and solve

$$\hat{\beta} = \arg \min_{\beta \in \mathcal{R}^p} \left\{ \frac{1}{n} \sum_{i=1}^n \frac{\phi(T_i X_i' \beta)}{\pi T_i + (1-T_i)/2} Y_i + \lambda_{1n} \|\beta\|_1 + \lambda_{2n} \eta(\beta) \right\}, \quad (4)$$

where $\|\beta\|_1$ is the L_1 norm of β . The second penalty $\eta(\beta)$ is optional and is imposed on the effects of ordinal variables with more than two levels. It could combine adjacent levels of such ordinal variables to result in simpler interpretation. In particular, we use a fused LASSO penalty (Tibshirani et al., 2004). For example, if one ordinal variable has C_j levels,

then the corresponding penalty term is $\sum_{m=2}^{C_j-1} |\beta_{j,m} - \beta_{j,m-1}|$, where $(\beta_{j,1}, \dots, \beta_{j,C_j-1}) \subset \beta$. If there are no ordinal variables or if there is no strong evidence that adjacent levels of an ordinal variable can have similar effects, we can also remove the corresponding fused LASSO penalty. As discussed in Section 1, the solution $\hat{\beta}$ directly reflects the covariate-treatment effects. Subgroups of patients that receive different treatment effects could be identified according to the non-zero elements of $\hat{\beta}$. Finally, our empirical rule is defined as $\hat{\mathcal{D}} = \text{sign}(X'\hat{\beta})$.

Since $\tilde{\beta}$ and β^* may be different, in Section 3 we propose an inference procedure to characterize the treatment effects for the identified subgroup. The validity of the inference

procedure hinges on the stability of $\tilde{\beta}$. We therefore prove a theoretical property of $\hat{\beta}$ to ensure the validity of the proposed procedure. We show in the Web Appendix A the consistency of $\hat{\beta}$ to $\tilde{\beta}$ as both n and p diverge to infinity with $n^{-1} \log p \rightarrow 0$. The results are shown for the case of bounded Y under the logistic loss: $\phi(x) = \log(1 + e^{-x})$, but can be easily extended to other types of outcomes. We also assume that $\mathbf{X}$ contains both categorical and continuous covariates as well as their pairwise interactions. In particular, we show that $\hat{\beta}$ possesses the weak oracle property (Fan and Lv, 2011), namely with large probability, $\hat{\beta}$ identifies all zero components of $\tilde{\beta}$ and gives consistent estimates to non-zero components of $\tilde{\beta}$ in a rate slower than $\sqrt{n}$.

3. Inference on Enhanced Comparative Treatment Effects

To evaluate the effect of subgroup identification, we consider the following quantity measuring the improvement of response when assigning subgroups of subjects to appropriate treatments:

$$\begin{aligned} d(\mathbf{X}, T, \hat{\beta}) &= E(Y | \hat{\mathcal{D}}(\mathbf{X}) = T) - E(Y | \hat{\mathcal{D}}(\mathbf{X}) \neq T) \\ &= E\{Y | \text{sign}(\mathbf{X}'\hat{\beta}) = T\} - E\{Y | \text{sign}(\mathbf{X}'\hat{\beta}) \neq T\} \end{aligned} \quad (5)$$

The actual comparative treatment effect for the subgroup identified for $T = 1$ via $\hat{\mathcal{D}}(\mathbf{X}) = 1$ is reflected by $d(\mathbf{X}, 1, \hat{\beta}) \equiv d(\mathbf{X}, T=1, \hat{\beta})$, and the mean effect is $E_{\mathbf{X}}\{d(\mathbf{X}, 1, \hat{\beta})\}$ for the identified subgroup. The expectation $E_{\mathbf{X}}$ is with respect to $\mathbf{X}$ and we treat β as “fixed” for future studies, similar to Foster et al. (2011). Correspondingly, the actual comparative treatment effect for the subjects identified for $T = -1$ is reflected by $d(\mathbf{X}, -1, \hat{\beta})$, and the mean effect is $E_{\mathbf{X}}\{d(\mathbf{X}, -1, \hat{\beta})\}$. Ideally, both $E_{\mathbf{X}}\{d(\mathbf{X}, 1, \hat{\beta})\}$ and $E_{\mathbf{X}}\{d(\mathbf{X}, -1, \hat{\beta})\}$ should be as large as possible. However it is also likely that one of them is small, for example when $T = -1$ is an inactive control. Inference about both $E_{\mathbf{X}}\{d(\mathbf{X}, 1, \hat{\beta})\}$ and $E_{\mathbf{X}}\{d(\mathbf{X}, -1, \hat{\beta})\}$ is important and can aid practical decision making regarding the identified subgroups.

In the following, we propose a bootstrap method to construct confidence intervals (CIs) for $E_{\mathbf{X}}\{d(\mathbf{X}, 1, \hat{\beta})\}$ and $E_{\mathbf{X}}\{d(\mathbf{X}, -1, \hat{\beta})\}$. If such intervals only contain positive numbers, we conclude that subgroup identification is indeed beneficial for treatment enhancement. We consider cases when $P\mathbf{X}'\beta^{\dagger} = 0 = 0$ and when $P\mathbf{X}'\beta^{\dagger} = 0 > 0$. The first case is regular and we can use the usual bootstrap. The second case is irregular (Shao, 1994; Laber and Murphy, 2011) due to a non-negligible probability concentrated at the discontinuous points of $d(\mathbf{X}, T, \hat{\beta})$ as a function of $\hat{\beta}$. In this case, we adopt the m -out-of- n bootstrap (Shao, 1994; Bickel, Götze, and van Zwet, 1997; Chakraborty, Laber, and Zhao, 2013). The method uses the bootstrap sample size m smaller than n to asymptotically avoid the non-negligible probability concentrated at the discontinuous points. Usually m is chosen as n^a for some $a < 1$. Chakraborty et al. (2013) used $a = 0.8$. We also adopt this choice in our simulation and

data analysis. Some preliminary sensitivity analysis on different choices of α is presented in the Web Appendix D. In practice, if many covariates are continuous, we can use the ordinary bootstrap (Moodie and Richardson, 2010). If many covariates are categorical, we propose to use m -out-of- n bootstrap.

Let $\{(Y_i, T_i, \mathbf{X}_i), i = 1, \dots, n\}$ be the original dataset. The m -out-of- n bootstrap algorithm comprises the following two steps.

Step 1: Generate a bootstrap sample with size m , $\{(Y_k^{(b)}, T_k^{(b)}, \mathbf{X}_k^{(b)}), k = 1, \dots, m\}$, from the original dataset. Obtain the corresponding estimates $\hat{\beta}^{(b)}$. Let $A_1 = \{k | \mathbf{X}_k' \hat{\beta}^{(b)} \geq 0, T_k = 1\}$, $A_2 = \{k | \mathbf{X}_k' \hat{\beta}^{(b)} \geq 0, T_k = -1\}$, $A_3 = \{k | \mathbf{X}_k' \hat{\beta}^{(b)} < 0, T_k = -1\}$, $A_4 = \{k | \mathbf{X}_k' \hat{\beta}^{(b)} < 0, T_k = 1\}$ and $|A_j|$ be the size of the set A_j for $j = 1, 2, 3$, and 4. Calculate $\tilde{d}(1, \hat{\beta}^{(b)}) = \sum_{k \in A_1} Y_k / |A_1| - \sum_{k \in A_2} Y_k / |A_2|$ and $\tilde{d}(-1, \hat{\beta}^{(b)}) = \sum_{k \in A_3} Y_k / |A_3| - \sum_{k \in A_4} Y_k / |A_4|$.

Step 2: Repeat *Step 1* B times. Let $\hat{l}_+$ and $\hat{u}_+$ be the $\alpha/2$ and $1-\alpha/2$ quantiles of $\{\tilde{d}(1, \hat{\beta}^{(b)}), b = 1, \dots, B\}$, and $\hat{l}_-$ and $\hat{u}_-$ corresponding quantiles of $\{\tilde{d}(-1, \hat{\beta}^{(b)}), b = 1, \dots, B\}$ respectively. Then, the $(1 - \alpha)$ CI of $E_{\mathbf{X}}\{d(\mathbf{X}, 1, \hat{\beta})\}$ is given by $[\hat{l}_+, \hat{u}_+]$ and the $(1 - \alpha)$ CI of $E_{\mathbf{X}}\{d(\mathbf{X}, -1, \hat{\beta})\}$ is given by $[\hat{l}_-, \hat{u}_-]$.

4. Numerical Results

In this section, we consider binary outcomes. Some of our simulation studies for time-to-event outcomes are presented in the Web Appendix B. We numerically compare our proposed regularized outcome weighted subgroup identification (ROWSi) method with some other recent subgrouping methods: (1) the FindIt by Imai and Ratkovic (in press), (2) the virtual twins (VT) by Foster et al. (2011), (3) the interaction tree (IT) by Su et al. (2009), (4) the logistic regression with LASSO penalty (logLASSO) by Qian and Murphy (2011), 5) the L_1 -penalized support vector machine (L1-SVM) method adapted from Zhao et al. (2012).

To implement the FindIt, we used the R package FindIt developed by Imai and Ratkovic with optimal tuning parameters chosen by their algorithm. We only considered two-way or three-way interactions among the treatment and covariates. To implement the VT, we first used the R function `randomForest` to estimate individual treatment effect $\hat{P}(y_i = 1 | T_i = 1, \mathbf{X}_i) - \hat{P}(Y_i = 1 | T_i = -1, \mathbf{X}_i)$ for subject i , and the average treatment effect $\hat{P}(y = 1 | T = 1) - \hat{P}(Y = 1 | T = -1)$. Then, we assigned subjects to treatment $T = 1$ if their individual treatment effects were greater than the average effect plus a constant, 0.05, as in Foster et al. (2011). We also followed Foster et al. (2011) for the choices of the terminal node size and complexity parameter. For the IT method, we first grew a large regression tree, where each split was induced by a threshold on some selected covariates. We then pruned the tree and selected the best-sized subtree based on an interaction-complexity

measurement and an additional penalty on the total number of internal nodes (Su et al., 2009).

In presence of a large number of interaction terms especially when many categorical variables exist, we adopted two options in our method. For the first option, we used a screening step similar to the sure independence screening method (Fan and Lv, 2008) to ease the computational burden of solving (4). In particular, we used a group LASSO procedure (Yuan and Lin, 2005) to pre-screen the interaction terms at the variable level in (4). For example, for the interaction between one variable with C_j levels and one variable with C_k levels, we treated the total $C_j \times C_k$ levels of interactions as a group and imposed a group LASSO penalty on it. For the second option, we abandoned this interaction screening step and directly solved (4). We chose tuning parameters λ_{1n} and λ_{2n} by the Akaike Information Criterion (AIC). We labeled our method using the latter option as “ROWSi_{onestep}” which allows us to investigate the advantage from the “screening step.”

All of our simulated data sets contained five binary covariates X_A, X_B, X_C, X_D, X_E , five ordinal covariates X_a, X_b, X_c, X_d, X_e and one continuous covariate X_{Ca} . Here we use capital letters in the subscript for binary variables, small letters for ordinal variables, and Ca for the continuous variable. For binary covariates, we use, for example X_{A1} , to denote the case level of X_A . For ordinal covariates, we use, for example X_{a3} to denote the third level of X_a . All binary covariates were simulated from the Bernoulli distribution with success probability of 1/2. All ordinal covariates had four uniformly distributed levels. The continuous covariate was simulated from the standard normal distribution. The treatment T was independent of the aforementioned covariates and took values of -1 and 1 with equal probabilities. We examined three different sample sizes: $n = 500, 1000$, and 2000 for each simulated data set and ran 200 simulations in each setting.

We considered a variety of models for logit $P(Y = 1)$ as follows.

- A.** $0.5X_{C1} + 2(X_{B1} + X_{a3} * X_{A1}) * T;$
- B.** $0.5X_{C1} + 2\{X_{B1} + X_{a3} * (X_{b2} + X_{b3})\} * T$
- C.** $0.05(-X_{A1} + X_{B1}) + \{(X_{a2} + X_{a3}) + (X_{b2} + X_{b3}) * X_{Ca}\} * T;$
- D.** $\log \log \{(X_{b3} + X_{c3}) + 5(X_{a2} + X_{a3} + X_{A1}X_{B1}) * T + 20\}^2;$
- E.** $(X_{A1} + X_{B1}) + 2 *;$
- F.** $0.5X_{A1} + 0.5X_{B1} + 2I\{X_{Ca} < 5, X_a < 2\} * T$
- G.** $0.5X_{A1} + 0.5X_{B1} + 2I\{X_{Ca} < 5, X_{Cb} < 2\} * T$
- H.** $0.5X_{Ca} + 0.5X_{Cb} + 2I\{X_{Ca} < -2, X_{Cb} > 2\} * T$

Models (A)–(C) are logistic models with various interaction structures. In particular, Model (A) has one main effect (X_{C1}), one first-order interactions ($X_{B1} * T$), and one second-order interaction ($X_{a3} * X_{A1} * T$). Model (B) is similar to (A) but has different second-order interactions. Model (C) has more complicated structure of main effects and interaction effects. Model (D) is not a logistic model and can be used to compare the robustness of different methods in case of link function misspecification. Model (E) is a logistic model

without covariate–treatment interactions. Models (F)–(H) take treelike forms which are not linear in terms of covariates.

We evaluated the performance of each method by four criteria: sensitivity, specificity, prediction accuracy, and the expected outcome $E^{\hat{\mathcal{D}}}(Y)$ under the resulting rules $\hat{\mathcal{D}}$ (a.k.a the value function, Zhao et al. (2012)). Sensitivity is the proportion of true non-zero interactions (including both first-order and second-order interactions) being estimated as non-zero. Specificity is the proportion of true zero interactions being estimated as zero. Prediction accuracy is the proportion of empirical treatment assignments, $\text{sign}(\mathbf{X}'\hat{\beta})$, that are consistent to the “oracle” assignments $\text{sign}(\mathbf{X}'\beta^*)$. Finally, interaction term selection accuracy is the proportion of a particular interaction term being selected.

To evaluate the sensitivity and specificity, Figure 1 plots the product of sensitivity and specificity for each method. Our method appears to outperform the competitors notably. Figure 2 plots prediction accuracy. Again, our method seemingly surpasses all other methods in almost all settings except under Model (H), where the true assignment rule is not linear. We present comparison of the value function in the Web Appendix C.

We used Models (A) and (C) to illustrate the construction of CI for enhanced comparative treatment effects. Model (A) does not satisfy the property that $P(\mathbf{X}'\beta^* = 0) = 0$ whereas Model (C) does. Hence, we used the m -out-of- n bootstrap for Model (A) and the regular bootstrap for Model (C). The bootstrap resampling was repeated for 1000 runs, from which a CI was obtained. We ran the above procedures for 1000 repetitions and report the mean length of CIs in Table 1. The length of CIs decreases as n increases. We also see that all CIs only include positive numbers, which indicates that treatment effects are enhanced in the subgroups identified by our method. In addition, we also report the coverage probabilities of the CIs for $E_{\mathbf{X}}\{d(\mathbf{X}, 1, \hat{\beta})\}$ and $E_{\mathbf{X}}\{d(\mathbf{X}, -1, \hat{\beta})\}$. Both of the two expectations were evaluated from a large independent testing data set with a sample size of 100,000. The performance seems satisfactory when the same size is 2000. There is some over-coverage when $n = 1000$. We noted that the performance depended on the choice of α in the m -out-of- n bootstrap. For this reason, we also investigated the sensitivity of this choice in Web Appendix D. Instead of a subjective choice, a double bootstrap procedure (Davison and Hinkley, 1997; Chakraborty et al., 2013) can be used for choosing α in a data-driven manner. Due to its computational intensity, we do not pursue here.

5. Data Analysis

In this section, we apply our proposed method to two real studies and compare it with the FindIt, VT and IT methods. The first is the National Supported Work (NSW) study (LaLonde, 1986) that appeared in Imai and Ratkovic (in press) and the FindIt package. The second is the mammography screening study discussed in the introduction section.

5.1. National Supported Work Study

In the NSW study, a training program was administered to a heterogeneous group of workers. The treatment is randomly assigned to each subject. It is of interest to investigate

whether the treatment effect varies as a function of individual characteristics. The treatment and control groups consist of 297 and 425 individuals, respectively. There are seven pre-analysis covariates including age, years of education, logarithm of annual earnings, race, marriage status, college degree status, and employment status. Including first and second order interactions, there are 45 covariates in total. The response is whether there is an increase on earnings from the year of 1975–1978.

Our method indicated that being married, having no high school degree and lower earnings had positive effects from the program; being Hispanics, employed, and having older age with higher earnings had negative effects. The FindIt method found that black people, people having lower education were positively affected by the program, but Hispanics who had higher earnings were negatively affected. The VT method found that the youth with lower earnings tended to get benefits. The IT method found that the white youth with lower earnings, or the youth with lower earnings and lower education tended to get benefits. The detailed resulting rules for subgroup identification from these methods are included in Web Appendix E.

To compare the performance of various methods, we randomly used 4/5 of samples (rounded to integers) as training sets to tune penalty parameters and the rest as test sets to evaluate the enhanced comparative treatment effects. The procedures were repeated for 200 random splits and the mean enhanced treatment effects are reported in Table 2. The estimated comparative treatment effects of each training set were estimated from the following formula:

$$\hat{E}_X\{d(\mathbf{X}, t, \hat{\beta})\} = \frac{\sum_{k=1}^{n_T} I(Y_k=1, \hat{\mathcal{D}}(\mathbf{X}_k)=T_k=t)}{\sum_{k=1}^{n_T} I(\hat{\mathcal{D}}(\mathbf{X}_k)=T_k=t)} - \frac{\sum_{k=1}^{n_T} I(Y_k=1, \hat{\mathcal{D}}(\mathbf{X}_k)=t, T_k=-t)}{\sum_{k=1}^{n_T} I(\hat{\mathcal{D}}(\mathbf{X}_k)=t, T_k=-t)},$$

where n_T is the sample size of the corresponding test set and $t = 1$ or -1 . The assignment rules $\hat{\mathcal{D}}(\mathbf{X}_k)$ were determined from the comparison methods and the training set. Results are shown in Table 2, including average subgroup sizes from the 200 random training sets. We see that our method led to higher increments of response rate than the other methods. Based on the whole data, we further calculated the 95% CIs for $E_X\{d(\mathbf{X}, 1, \hat{\beta})\}$ and $E_X\{d(\mathbf{X}, -1, \hat{\beta})\}$ using the m -out-of- n bootstrap, which were (6.92%, 23.90%) and (3.03%, 21.41%), respectively. Since 0 is not contained in the above two intervals, we conclude that the treatment effects were likely enhanced in the subgroups identified by our method.

5.2. Mammography Screening Study

This is a randomized study that included female subjects who were non-adherent to mammography screening guidelines at baseline (i.e., no mammogram in the year prior to baseline) (Champion et al., 2007). One primary interest of the study was to compare the intervention effects of phone counseling on mammography screening (phone intervention)

versus usual care at 21 months post-baseline. The outcome is whether the subject took mammography screening during this time period. There are 530 subjects with 259 in the phone intervention group and 271 in the usual care group. There are 16 binary variables, including socio-demographics, health belief variables, and stage of readiness to undertake mammography screening, and 1 categorical variable, number of years had a mammogram in past 2 to 5 years, in the study. Considering the covariates' main effects and the first and second order interactions with T , there are 409 covariates in total.

Our method found that people having more past mammograms or with higher self efficacy score benefited from the phone intervention. On the other hand, people with low household income or higher barriers scale score were negatively affected by the phone intervention. The FindIt method did not find any covariate–treatment interaction. The VT method found that people having more past mammograms and at the ready stage with higher household income, or married people with fewer past mammograms and at the ready stage benefited from the phone intervention. The IT method found that people having more past mammograms, or people who were at the ready stage with higher household income benefited from the phone intervention. The resulting rules from these methods are included in Web Appendix E.

Similar to the cross-validation procedure we used for Table 2, we report the average estimated enhanced treatment effects and subgroup sizes in Table 3. Due to the smaller sample size, we used half-half random splits. Our identified subgroups again have higher increments of response rate than the other methods on average. We were not able to obtain $\hat{E}_X\{d(\mathbf{X}, -1, \hat{\beta})\}$ for FindIt because it did not find any covariate-treatment interaction in many training sets. We further calculated the 95% CIs for $E_X\{d(\mathbf{X}, 1, \hat{\beta})\}$ and $E_X\{d(\mathbf{X}, -1, \hat{\beta})\}$ for our method using the m -out-of- n bootstrap, which were (5.88%, 11.21%) and (4.39%, 10.91%) respectively. The CIs do not include 0, indicating plausible identifications for subgroups.

6. Conclusion and Discussion

In this article, we proposed a regularized outcome weighted method for subgroup identification. The method differs from many existing methods for subgroup identification in that it is based on a target function that directly relates to covariate-treatment interaction. Such separation from the covariate main effects is important as these main effects usually affect the estimation of interactions. Because the clinical outcomes are used as weights instead of as modeling targets, our method can be extended to deal with continuous, time-to-event, and possibly contaminated outcomes. Another advantage of our method is that it could be easily implemented by using standard packages. In fact, under the logistic surrogate loss $\phi(x) = \log(1 + e^{-x})$, if the regularization part of (4) is ignored, we can essentially solve

$$\hat{\beta} = \arg \min_{\beta \in \mathbb{R}^p} \left\{ \frac{1}{n} \sum_{i=1}^n \frac{\log(1 + e^{-T_i \mathbf{X}_i' \beta})}{\pi T_i + (1 - T_i)/2} Y_i \right\}.$$

based on a logistic regression model for T weighted by outcomes. This can be conveniently fitted using standard packages such as the SAS LOGISTIC procedure by using subject-specific weights. For more general settings, we have presented some codes in the Web Appendix F. To keep it easier for the readers to implement, we have only presented codes that use LASSO as the penalty. We are developing an R package that will also incorporate the fused LASSO penalty.

We demonstrated the impressive performance of our method compared with other methods in various settings. Moreover, we not only identified the subgroups, but also carried out an additional inference procedure to assess how much benefit such a subgroup identification procedure could bring. Our approach for subgroup identification is mainly used to determine the rules for treatment assignment. This should be distinguished from other approaches for “subgroup identification” that focus on identifying some natural regions in the covariate space with enhanced treatment effects with no implication that the complement of that region constitutes another subgroup with a negative treatment effect (Kehl and Ulm, 2006; Lipkovich et al., 2011; Dusseldorp and Van Mechelen, 2014). This distinction is important as it may lead to different strategies in personalized medicine: one is to find the right drugs for patients and the other to find right patients for drugs.

For simplicity of presentation, we have assumed a complete randomization scheme from clinical trial data, that is $P(T=1) = \pi$. Our method also works when the treatment assignment depends on X , that is, $P(T=1|X) = \pi(X)$ by replacing $T\pi + (1-T)/2$ in (1)–(4) with $T\pi(X) + (1-T)/2$. The approach can also be extended to multiple interval treatments. The framework will be similar to dynamic treatment regimes, which are customized sequential decision rules for individual patients that can adapt over time to an evolving illness (Murphy, 2003; Zhao et al., in press). We can generalize the proposed method based on dynamic programming, which identifies the best rule at each interval backwards and recursively.

Zhang et al. (2012) showed that the optimal treatment regime could be regarded as the solution of the classification problem that

$\mathcal{D}^* = \arg\min_{\mathcal{D}} E(|C(X)|[I\{C(X) > 0\} - \mathcal{D}(X)]^2)$, in which $C(X) = E(Y|T=1, X) - E(Y|T=0, X)$ is the treatment contrast. Our formulation in (1) can be generalized by using $C(X)$ based subject-specific weight. However, $C(X)$ is unknown and needs to be estimated. Zhang et al. (2012) proposed a doubly robust estimator of $C(X)$, which protects the estimation against model misspecification, especially when treatment assignment probabilities are known as in a randomized clinical trial. It will be of great interest to further investigate such a procedure as well as its theoretical properties.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

The research efforts were partly supported by the University of Wisconsin Carbone Cancer Center support grant, NIH/NIA P30 CA014520 (for Drs. Menggang Yu and Yaoyao Xu), a development fund from the University of Wisconsin Carbone Cancer Center (for Dr. Menggang Yu), and NIH 5 R01 HG007377-02 (for Dr. Sijian Wang).

References

- Bartlett PL, Jordan MI, McAuliffe JD. Convexity, classification, and risk bounds. *Journal of the American Statistical Association*. 2006; 101:138–156.
- Bickel PJ, Götze F, van Zwet WR. Resampling fewer than n observations: Gains, losses, and remedies for losses. *Statistica Sinica*. 1997; 7:1–31.
- Breiman, L., Friedman, JR., Stone, CJ., Olshen, R. *Classification and Regression Trees*. Wadsworth: Chapman and Hall; 1984.
- Brookes ST, Whitely E, Egger M, Smith GD, Mulheran PA, Peters TJ, et al. Subgroup analyses in randomized trials: Risks of subgroup-specific analyses; power and sample size for the interaction test. *Journal of Clinical Epidemiology*. 2004; 57:229. [PubMed: 15066682]
- Cai T, Tian L, Wong PH, Wei L. Analysis of randomized comparative clinical trial data for personalized treatment selections. *Biostatistics*. 2011; 12:270–282. [PubMed: 20876663]
- Chakraborty B, Laber EB, Zhao Y. Inference for optimal dynamic treatment regimes using an adaptive m-out-of-n bootstrap scheme. *Biometrics*. 2013; 69:714–23. [PubMed: 23845276]
- Champion V, Skinner CS, Hui S, Monahan P, Juliar B, Daggy J, Menon U. The effect of telephone v. print tailoring for mammography adherence. *Patient Education and Counseling*. 2007; 65:416. [PubMed: 17196358]
- Davison, AC., Hinkley, DV. *Bootstrap Methods and Their Applications*. Cambridge: Cambridge University Press; 1997.
- Dusseldorp E, Van Mechelen I. Qualitative interaction trees: A tool to identify qualitative treatment Subgroup interactions. *Statistics in Medicine*. 2014; 33:219–237. [PubMed: 23922224]
- Fan J, Lv J. Sure independence screening for ultrahigh dimensional feature space. *Journal of the Royal Statistical Society, Series B*. 2008; 70:849–911.
- Fan J, Lv J. Nonconcave penalized likelihood with npdimensionality. *IEEE Transactions on Information Theory*. 2011; 57:5467–5484. [PubMed: 22287795]
- Foster JC, Taylor JM, Ruberg SJ. Subgroup identification from randomized clinical trial data. *Statistics in Medicine*. 2011; 30:2867–2880. [PubMed: 21815180]
- He, X. PhD thesis. University of Wisconsin-Madison; 2012. Identification of subgroups with large differential treatment effects in genome-wide association studies.
- Imai K, Ratkovic M. Estimating treatment effect heterogeneity in randomized program evaluation. *Annals of Applied Statistics*. 2013; 7:443–470.
- Janes H, Pepe M, Bossuyt P, Barlow W. Measuring the performance of markers for guiding treatment decisions. *Annals of Internal Medicine*. 2011; 154:253–9. [PubMed: 21320940]
- Janz N, Becker M. The health belief model: A decade later. *Health Education Quarterly*. 1984; 11:1–47. [PubMed: 6392204]
- Kehl V, Ulm K. Responder identification in clinical trials with censored data. *Computational Statistics & Data Analysis*. 2006; 50:1338–1355.
- Laber EB, Murphy SA. Adaptive confidence intervals for the test error in classification. *Journal of the American Statistical Association*. 2011; 106:904–13. [PubMed: 22053123]
- Lagakos SW. The challenge of subgroup analyses – Reporting without distorting. *New England Journal of Medicine*. 2006; 354:1667–1669. [PubMed: 16625007]
- LaLonde R. Evaluating the econometric evaluations of training programs with experimental data. *American Economic Review*. 1986; 76:604–20.
- Li KC. Sliced inverse regression for dimension reduction. *Journal of American Statistical Association*. 1991; 86:316–27.

- Lipkovich I, Dmitrienko A, Denne J, Enas G. Subgroup identification based on differential effect search: A recursive partitioning method for establishing response to treatment in patient subpopulations. *Statistics in Medicine*. 2011; 30:2601–2621. [PubMed: 21786278]
- Moodie EE, Richardson TS. Bias correction in non-differentiable estimating equations for optimal dynamic regimes. *Scandinavian Journal of Statistics*. 2010; 37:126–146.
- Murphy S. Optimal dynamic treatment regimes (with discussion). *Journal of the Royal Statistical Society, Series B*. 2003; 65:331–66.
- Qian M, Murphy SA. Performance guarantees for individualized treatment rules. *Annals of Statistics*. 2011; 39:1180. [PubMed: 21666835]
- Ruberg SJ, Chen L, Wang Y. The mean does not mean as much anymore: Finding sub-groups for tailored therapeutics. *Clinical Trials*. 2010; 7:574–583. [PubMed: 20667935]
- Sargent D, Mandrekar S. Statistical issues in the validation of prognostic, predictive, and surrogate biomarkers. *Clinical Trials*. 2013; 10:647–52. [PubMed: 23983158]
- Shao J. Bootstrap sample size in nonregular cases. *Proceedings of the American Mathematical Society*. 1994; 122:1251–1262.
- Song R, Wang W, Zeng D, Kosorok M. Penalized q-learning for dynamic treatment regimens. *Statistica Sinica*. 2015; 25
- Su X, Tsai CL, Wang H, Nickerson DM, Li B. Subgroup analysis via recursive partitioning. *The Journal of Machine Learning Research*. 2009; 10:141–158.
- Su X, Zhou T, Yan X, Fan J, Yang S. Interaction trees with censored survival data. *The International Journal of Biostatistics*. 2008; 4:1–26.
- Tibshirani R, Saunders M, Rosset S, Zhu J, Knight K. Sparsity and smoothness via the fused lasso. *Journal of the Royal Statistical Society, Series B*. 2004; 67:91–108.
- Varadhan R, Segal JB, Boyd CM, Wu AW, Weiss CO. A framework for the analysis of heterogeneity of treatment effect in patient-centered outcomes research. *Journal of Clinical Epidemiology*. 2013; 66:818–825. [PubMed: 23651763]
- Yuan M, Lin Y. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society, Series B*. 2005; 68:49–67.
- Zhang B, Tsiatis A, Davidian M, Zhang M, Laber E. Estimating optimal treatment regimes from a classification perspective. *Stat*. 2012; 1:103–114. [PubMed: 23645940]
- Zhao Y, Zeng D, Laber E, Kosorok M. New statistical learning methods for estimating optimal dynamic treatment regimes. *Journal of the American Statistical Association*. (in press).
- Zhao Y, Zeng D, Rush A, Kosorok M. Estimating individualized treatment rules using outcome weighted learning. *Journal of the American Statistical Association*. 2012; 107:1106–1118. [PubMed: 23630406]

Figure 1. Comparison of the product of sensitivity and specificity by seven methods, represented by various types of points and lines as in the legend of Model A (number of simulations = 200).

Figure 2. Comparison of prediction accuracy by seven methods, represented by various types of points and lines as in the legend of Model A (number of simulations = 200).

Confidence intervals of enhanced comparative treatment effects in Models (A) and (C). Reported values are the means from 1000 repetitions, each of which uses 1000 subsamples for the bootstrap. For notational simplicity, we denote $d_+ \equiv E_{\mathbf{X}}\{d(\mathbf{X}, 1, \hat{\beta})\}$ and $d_- \equiv E_{\mathbf{X}}\{d(\mathbf{X}, -1, \hat{\beta})\}$.

Table 1

	$n = 500$		$n = 1000$		$n = 2000$	
	d_+	d_-	d_+	d_-	d_+	d_-
Model (A)						
CI length	0.55	0.57	0.41	0.34	0.24	0.24
coverage	0.93	0.98	0.97	0.99	0.96	0.96
Model (C)						
CI length	0.17	0.19	0.13	0.14	0.09	0.10
coverage	0.93	0.94	0.97	0.97	0.95	0.96

Table 2

NSW study: the average estimated enhanced comparative treatment effects and subgroup sizes based on cross-validation. The cross-validation used 200 random splits and reported values are the means from these random splits. For notational simplicity, we denote $\hat{d}_+(\hat{\beta}) \equiv \hat{E}_{\mathbf{X}}\{d(\mathbf{X}, 1, \hat{\beta})\}$ and $\hat{d}_-(\hat{\beta}) \equiv \hat{E}_{\mathbf{X}}\{d(\mathbf{X}, -1, \hat{\beta})\}$

	Enhanced treatment effect (subgroup size)			
	ROWSi	FindIt	VT	IT
$\hat{d}_+(\hat{\beta})$	18.9% (121)	5.61% (74)	3.53% (97)	3.66% (106)
$\hat{d}_-(\hat{\beta})$	4.5% (23)	-2.34% (70)	-3.84% (47)	-5.66% (38)

Table 3

Mammography screening study: the average estimated enhanced comparative treatment effects and subgroup sizes based on cross-validation. The cross-validation used 200 random splits and reported values are the means from these random splits. For notational simplicity, we denote $\hat{d}_+(\hat{\beta}) \equiv \hat{E}_{\mathbf{X}}\{d(\mathbf{X}, 1, \hat{\beta})\}$ and

$$\hat{d}_-(\hat{\beta}) \equiv \hat{E}_{\mathbf{X}}\{d(\mathbf{X}, -1, \hat{\beta})\}$$

	Enhanced treatment effect (subgroup size)			
	ROWSi	FindIt	VT	IT
$\hat{d}_+(\hat{\beta})$	7.88% (149)	2.83% (265)	3.58% (135)	3.42% (126)
$\hat{d}_-(\hat{\beta})$	6.89% (116)	NA	0.38% (130)	3.10% (139)