Published in final edited form as:

Biometrics. 2014 December ; 70(4): 891–901. doi:10.1111/biom.12204.

Identifying Optimal Biomarker Combinations for Treatment Selection via a Robust Kernel Method

Ying Huang* and Youyi Fong

Fred Hutchinson Cancer Research Center Public Health Sciences, 1100 Fairview Avenue N., Seattle, WA 98109-1024, University of Washington, Department of Biostatistics, Seattle, WA, 98195

Summary

Treatment-selection markers predict an individual's response to different therapies, thus allowing for the selection of a therapy with the best predicted outcome. A good marker-based treatment-selection rule can significantly impact public health through the reduction of the disease burden in a cost-effective manner. Our goal in this paper is to use data from randomized trials to identify optimal linear and nonlinear biomarker combinations for treatment selection that minimize the total burden to the population caused by either the targeted disease or its treatment. We frame this objective into a general problem of minimizing a weighted sum of 0–1 loss and propose a novel penalized minimization method that is based on the difference of convex functions algorithm (DCA). The corresponding estimator of marker combinations has a kernel property that allows flexible modeling of linear and nonlinear marker combinations. We compare the proposed methods with existing methods for optimizing treatment regimens such as the logistic regression model and the weighted support vector machine. Performances of different weight functions are also investigated. The application of the proposed method is illustrated using a real example from an HIV vaccine trial: we search for a combination of Fc receptor genes for recommending vaccination in preventing HIV infection.

Keywords

Biomarker Combination; Kernel Method; Randomized Trial; Robust; Support Vector Machine; Treatment Selection

1. Introduction

Heterogeneity exists among individual subjects' responses to treatment in many disease settings. Biomarkers associated with treatment-response heterogeneity can help individuals select therapies in order to optimize clinical outcome. As a single marker typically has limited ability to predict the heterogeneity in treatment response, statistical methods for combining biomarkers are critically needed for optimizing treatment selection.

*yhuang@fhcrc.org.

Among others, one common approach for identifying marker combinations in treatment selection relies on modeling of the disease risk, conditional on biomarkers and treatment assignment; subsequent treatment recommendation is made based on the difference in the disease risk between non-treated and treated. For example, Song and Pepe (2004) adopted the standard logistic regression for disease risk modeling; Foster et al. (2011) proposed a ‘virtual twin’ method with tree-based estimators of disease risk; Qian and Murphy (2011) and McKeague and Qian (2013) adopted a L1-penalized least squares approach for estimating treatment-selection rules; Lu et al. (2013) developed an A-learning method for linear outcomes that requires correct modeling of the treatment-marker interaction but is robust to modeling of the marker’s main effect.

Performance of the treatment-selection rules based on risk modeling relies critically on the correct specification of the model, which is challenging given the complexity of biological mechanisms. A different approach that is more robust to model misspecification is to write down a desired criterion function associated with treatment selection, and search for the marker combination that optimizes the criterion function. For example, Orellana et al. (2010) considered parametric and semiparametric dynamic regime marginal structural models for optimizing expected utility. To optimize the population mean outcome, Zhang et al. (2012a, b) proposed finding marker combinations that optimize the empirical or double-robust estimates of population mean outcome within a pre-specified parametric class. Since these estimates of population mean outcome involve a non-smooth, non-convex 0–1 loss, the optimization procedure can be computationally challenging. Zhao et al. (2012) transformed the optimization problem into an outcome-weighted learning problem and derived the optimal treatment-selection rule using a weighted support vector machine method by approximating the 0–1 loss with a convex hinge loss function.

In this paper, we propose a new method to identify marker combinations that directly optimize the targeted criterion function as did Zhang et al. (2012a) and Zhao et al. (2012). In particular, we consider the objective of minimizing unfavorable outcomes in the population. These outcomes possibly include 1) events of a targeted disease, and 2) adverse side effects or monetary cost incurred by treatment. Furthermore, in a randomized trial setting, the optimization problem can be reformulated as one of minimizing a weighted sum of 0–1 loss. We propose to solve this via the difference of convex functions algorithm (DCA). Like the outcome-weighted support vector machine in Zhao et al. (2012), our estimator allows for flexible linear and nonlinear combinations of biomarkers to account for the possibly complicated relationship between biomarkers and treatment. However, our estimator is more effective at minimizing the targeted criterion function through the adoption of a more precise approximation of 0–1 loss. Fisher consistency and asymptotic consistency of the proposed estimator are derived. Different choices of weight functions are also investigated through numerical studies.

Next, in Section 2, we describe a general framework for optimizing treatment-selection rules. In Section 3, we propose a penalized estimator that minimizes the penalized weighted sum of a ramp loss function and develop its theoretical properties. In Section 4, we evaluate the performance of the estimator and compare it with common approaches through

numerical studies. In Section 5, we illustrate the methodology with an example from a recent HIV vaccine efficacy trial in Thailand. Finally, we end the paper with a discussion.

2. Method

Here we consider the problem of optimizing the treatment regimen for an unfavorable clinical outcome denoted by Y , based on a set of markers $\mathbf{X}$ with dimension $p \geq 1$. In this paper we focus on a binary disease outcome Y , 0 for non-diseased and 1 for diseased. The methods developed work for continuous outcomes as well. Let $A(\mathbf{X})$ indicate the marker-based treatment-selection rule. $A = 0$ stands for not treating and $A = 1$ for treating. An important measure for quantifying the treatment-selection benefit of $A(\mathbf{X})$ is $E_{A(\mathbf{X})}(Y)$, the expected disease rate as a result of treatment selection (Song and Pepe, 2004). Here the expectation is averaged over the decision $A(\mathbf{X})$. This measure has become more widely recognized in the literature in recent years (Zhao et al., 2012; Zhang et al., 2012a).

While $E_{A(\mathbf{X})}(Y)$ characterizes the burden of disease upon the population, more generally, one may be concerned with additional burden associated with treatment. In practice there are often unpleasant aspects associated with treatment such as its side effects and/or monetary cost. A treatment-selection strategy that takes these aspects into consideration is thus preferred. One way to combine the disease and treatment burdens is to pre-specify a treatment/disease harm ratio such that each burden type can be put on the same scale. As in the decision-theoretic framework proposed in Vickers et al. (2007), let δ be a pre-specified ratio of the burden per treatment relative to the burden per disease event, and let $Y(1)$ and $Y(0)$ indicate the potential disease outcome if a subject were to receive or not receive the treatment. Here we set the burden per targeted disease event as 1. Then the total burden due to disease and treatment for a treatment-selection rule $A(\mathbf{X})$, represented in the unit of the burden per disease event, equals

$$\theta = E_{A(\mathbf{X})}(Y) + E\{\delta \times A(\mathbf{X})\} = \sum_{a=0}^1 E[Y(a) \times I\{A(\mathbf{X})=a\}] + \delta \times E\{A(\mathbf{X})\}. \quad (1)$$

An optimal treatment-selection rule can thus be found by minimizing θ .

In practice, it is often difficult to agree upon δ . Choosing $\delta = 0$ such that $\theta = E_{A(\mathbf{X})}(Y)$ corresponds to the special case where the criterion function to be minimized is the burden due to the targeted disease alone, as in Zhang et al. (2012a) and Zhao et al. (2012). While detailed discussion about the choice of δ in treatment selection is beyond the scope of this paper, the development of a generalized framework allowing for nonzero δ is, however, important for a method to be applicable to practical scenarios where there are sensible ways to pre-specify δ . For example, δ was chosen to be 5% for treating breast cancer with adjuvant chemotherapy in Vickers et al. (2007) according to a patient survey. In Rapsomaniki et al. (2012), δ was chosen to be 2% for cost of preventive treatment per year for cardiovascular disease relative to the value of an event-free life year.

Suppose we have data from a two-arm randomized trial with $T = 0, 1$ indicating assignment to the untreated and treated arm, respectively. Let n_0 and n_1 indicate the number of subjects

untreated and treated respectively. We have i.i.d. samples $\{Y_i, \mathbf{X}_i, T_i\}$ for $i = 1, \dots, n$ with $n = n_0 + n_1$. We make the following assumptions: (i) stable unit treatment value (SUTVA) (Rubin, 1980) and consistency: $Y(0), Y(1)$ of one subject is independent of the treatment assignments of other subjects, and given the treatment a subject actually received, a subject's potential outcomes equal the observed outcomes; (ii) ignorable treatment assignments assumption: $T \perp Y(0), Y(1) | \mathbf{X}$. Assumption (i) is plausible in trials where participants do not interact with one another and assumption (ii) is ensured by randomization. Algebra (as in Web Supplementary Appendix A) shows that

$$\theta = E[Y \times \{1 - A(\mathbf{X})\} | T=0] + E[Y \times A(\mathbf{X}) | T=1] + \delta \times P\{A(\mathbf{X})=1\} \quad (2)$$

$$= E(Y | T=0) - E[A(\mathbf{X}) \times \{Risk_0(\mathbf{X}) - Risk_1(\mathbf{X}) - \delta\}], \quad (3)$$

where $Risk_0(\mathbf{X}) = P(Y = 1 | \mathbf{X}, T = 0)$ and $Risk_1(\mathbf{X}) = P(Y = 1 | \mathbf{X}, T = 1)$ are the risk of Y conditional on $\mathbf{X}$ among untreated and treated, respectively. Therefore an optimal rule $A(\mathbf{X})$ in the sense of minimizing θ is $A(\mathbf{X}) = 1$ if $Risk_0(\mathbf{X}) - Risk_1(\mathbf{X}) > \delta$ and $A(\mathbf{X}) = 0$ otherwise. In other words, the reduction in the disease burden from treatment needs to be more significant than the burden of treatment itself for one to recommend it for a subject. Note that $E(Y | T=0) - \theta$ is the net benefit of $A(\mathbf{X})$ as defined in Vickers et al. (2007), i.e., the difference in the total burden comparing treating none with $A(\mathbf{X})$. It equals $E[A(\mathbf{X}) \times \{Risk_0(\mathbf{X}) - Risk_1(\mathbf{X}) - \delta\}]$ by equation (3) and is maximized by the optimal rule $A(\mathbf{X}) = I\{Risk_0(\mathbf{X}) - Risk_1(\mathbf{X}) > \delta\}$. Related results have been shown by others. For example, in the setting where Y is defined to be a favorable outcome, Zhang et al. (2012b) showed that the optimal $A(\mathbf{X})$ when $\delta = 0$ is to treat if $Risk_1(\mathbf{X}) - Risk_0(\mathbf{X}) > 0$; Baker et al. (2012) showed that the average treatment benefit in Y needs to be larger or smaller than δ among subjects with $A(\mathbf{X}) = 1$ or 0 respectively, in order for $A(\mathbf{X})$ to be more beneficial than treating none and treating all: $A(\mathbf{X}) = I\{Risk_1(\mathbf{X}) - Risk_0(\mathbf{X}) > \delta\}$ satisfies these conditions. Standard regression-based methods can be used to estimate $Risk_0(\mathbf{X})$ and $Risk_1(\mathbf{X})$ in order to derive the optimal treatment-selection rule. As pointed out by Zhang et al. (2012a), however, when the risk model is misspecified, performance of the treatment-selection rule derived this way can be sub-optimal.

To make treatment selection more robust to model misspecification, we follow the strategy of Zhang et al. (2012a) and Zhao et al. (2012), and consider a class of rules that make treatment recommendation based on whether a marker combination score is greater than zero. In particular, let $A(\mathbf{X}) = I\{f(\mathbf{X}) > 0\}$ with I the indicator function, $f(\mathbf{X}) = b + \gamma(\mathbf{X})$ with $\gamma(\mathbf{X})$ a function of markers $\mathbf{X}$. Then θ in (2) can equivalently be represented as

$$\theta = \frac{E\{Y \times T \times I\{f(\mathbf{X}) > 0\}\}}{P(T=1)} + \frac{E\{Y \times (1-T) \times I\{f(\mathbf{X}) \leq 0\}\}}{P(T=0)} + \delta \times P\{f(\mathbf{X}) > 0\}.$$

The optimal $f(\mathbf{X})$ can be found by minimizing the empirical estimate of θ , i.e.,

$$\begin{aligned}\hat{f} &= \operatorname{argmin}_f \sum_{i=1}^n \left[\frac{Y_i \times T_i \times I\{f(\mathbf{X}_i) > 0\}}{n_1} + \frac{Y_i \times (1-T_i) \times I\{f(\mathbf{X}_i) \leq 0\}}{n_0} + \delta \times \frac{I\{f(\mathbf{X}_i) > 0\}}{n} \right] \\ &= \operatorname{argmin}_f \sum_{i=1}^n -\{Y_i \times T_i/n_1 - Y_i \times (1-T_i)/n_0 + \delta/n\} \times I\{f(\mathbf{X}_i) \leq 0\}.\end{aligned}$$

Therefore, we can formulate this problem as the minimization of a weighted sum of 0–1 loss. That is, $\hat{f}$ can be found as the minimizer of

$$\sum_{i=1}^n W_i I\{f(\mathbf{X}_i) \leq 0\}/n \propto \sum_{i=1}^n |W_i| I\{\operatorname{sgn}(W_i) \times f(\mathbf{X}_i) \leq 0\}/n \quad (4)$$

with the subject-specific weight

$$W_i = W_{1i} = -\left\{ \frac{Y_i \times T_i}{n_1/n} - \frac{Y_i \times (1-T_i)}{n_0/n} + \delta \right\}. \quad (5)$$

Note that in the special case where $\delta = 0$, we have $W_{1i} = 0$ when $Y_i = 0$, which implies that we only need to measure biomarkers among cases to derive the optimal treatment-selection rule. In subsequent numerical studies, we name this set of weights the ‘case-only weights’.

Second, note that θ can also be represented as $E(Y|T=0) + E[(1-Y) \times (1-T) \times I\{f(\mathbf{X}) > 0\}]/P(T=0) - E[(1-Y) \times T \times I\{f(\mathbf{X}) > 0\}]/P(T=1) + \delta \times P\{f(\mathbf{X}) > 0\}$. So to estimate $f(\mathbf{X})$

we can minimize $\sum_{i=1}^n \{(1-Y_i) \times (1-T_i)/n_0 - (1-Y_i) \times T_i/n_1 + \delta/n\} \times I\{f(\mathbf{X}_i) > 0\}$, which is equivalent to minimizing the weighted sum of 0–1 loss (4) with the weight

$$W_i = W_{2i} = -\left\{ \frac{(1-Y_i) \times (1-T_i)}{n_0/n} - \frac{(1-Y_i) \times T_i}{n_1/n} + \delta \right\}. \quad (6)$$

Note that for $\delta = 0$, using this set of weights only requires collection of biomarker information among controls since $W_{2i} = 0$ for $Y_i = 1$. In subsequent numerical studies, we name this set of weights the ‘control-only weights’.

Since the minimizer of two different sets of loss will also be the minimizer of the linear combination of the two sets of loss, we can choose weights in loss (4) as a weighted average of the case-only weights (5) and the control-only weights (6) with weight p and $1-p$ respectively for arbitrary p . This is equivalent to minimizing (4) with

$$W_i = W_{3i} = -\left[\{p \times Y_i - (1-p) \times (1-Y_i)\} \times \left(\frac{T_i}{n_1/n} - \frac{1-T_i}{n_0/n} \right) + \delta \right]. \quad (7)$$

In subsequent numerical studies, we choose $p = 1/2$ and name the set of weights the ‘case-control weights’. All three sets of weights W_1 , W_2 , W_3 do not require modeling of the disease risk conditional on the marker and the treatment.

Finally, we consider substituting the doubly robust estimator of $E_{A(X)}(Y)$ (Zhang et al., 2012a) for estimating θ in (2). This estimator augments the inverse probability weighted estimator of $E_{A(X)}(Y)$ with a term that involves the risk of Y given X and T . It has the double robustness property in that it is consistent for $E_{A(X)}(Y)$ if either $P(T = 1|X)$ (the propensity score for treatment) or the risk model is correctly specified. In a randomized trial, the propensity score for treatment is known, so consistency of the estimator is always achievable. Using the augmented term is expected to have the added advantage of improving efficiency, as shown in Zhang et al. (2012a) for continuous outcomes. Using a working model to obtain risk estimates $\hat{Risk}_0(X)$ and $\hat{Risk}_1(X)$ leads to the estimate of θ as

$$\hat{\theta} = \sum_{i=1}^n \left[\frac{Y_i \times (1-T_i) \times I\{f(X_i) \leq 0\}}{n_0} + \frac{Y_i \times T_i \times I\{f(X_i) > 0\}}{n_1} + \delta \times \frac{I\{f(X_i) > 0\}}{n} \right] - \frac{\sum_{i=1}^n ((1-T_i) \times I\{f(X_i) < 0\} + T_i \times I\{f(X_i) > 0\} - P[I\{f(X_i) > 0\} = T|X_i]) \times \hat{m}(Y_i; f)}{P[I\{f(X_i) > 0\} = T|X_i]},$$

where $P[I\{f(X) > 0\} = T|X] = \pi \times I\{f(X) > 0\} + (1 - \pi) \times I\{f(X) \leq 0\}$ for $\pi = P(T = 1)$, $\hat{m}(X_i; f) = \hat{Risk}_0(X_i) \times I\{f(X_i) \leq 0\} + \hat{Risk}_1(X_i) \times I\{f(X_i) > 0\}$. As shown in Web Supplementary Appendix B, this corresponds to minimization of (4) with the weight

$$W_i = W_{4i} = - \left[\frac{Y_i T_i}{n_1} - \frac{Y_i (1-T_i)}{n_0} + \frac{\pi - T_i}{1-\pi} \times \frac{\hat{Risk}_0 + \hat{Risk}_1}{n} + \frac{\delta}{n} \right]. \quad (8)$$

We name this set of weights the ‘double-robustness weights’.

In subsequent numerical studies, we evaluate the performance of our estimators based on the above four special weights. We note, however, the existence of alternative ways to generate the weights. In particular, since the minimization of (3) is equivalent to the minimization of $E[I\{f(X) \leq 0\} \times \{Risk_0(X) - Risk_1(X) - \delta\}]$, any consistent estimate of $Risk_0(X) - Risk_1(X) - \delta$ can serve as W in (4). Similar observation has been made for the problem of minimizing $E_{A(X)}(Y)$ (Zhang et al., 2012b).

2.1 Penalized Weighted Ramp Loss Estimator

In this section we consider the minimization of the loss function (4) conditional on a pre-specified set of weights W . Minimization of a weighted sum of 0–1 loss is computationally challenging due to the non-convexity of the indicator function. We propose to derive the marker combination through a difference of convex functions algorithm (Liu et al., 2005; Wu and Liu, 2007). Specifically, we approximate the indicator function $I(u \leq 0)$ with a ramp loss function $h_s(u)$ as shown in Figure 1(a), where $h_s(u) = 1$ if $u/s \leq -1/2$; $h_s(u) = -u/s + 1/2$ if $-1/2 \leq u/s \leq 1/2$; and $h_s(u) = 0$ if $u/s \geq 1/2$, with $s > 0$ a scale factor. For any given $u \leq 0$, the value of this loss function moves towards 0 or 1 and becomes 0 or 1 at some point as $s \rightarrow 0$. Identifiability issues exist in minimizing the weighted sum of ramp loss because $h_s(|u|)$ is greater than or equal to $h_s(c|u|)$ for $c > 1$ and the ramp loss function satisfies $h_s(|u|) = h_{s=1}(|u|/s)$. We handle the identifiability problem and at the same time control the model complexity by setting $s = 1$ and adding a penalty term to the weighted ramp loss. Using $h(u)$

to indicate $h_s(u)$ with $s = 1$, we propose to estimate $f(\mathbf{X}) = b + \gamma(\mathbf{X})$ by minimizing a penalized weighted sum of ramp loss

$$\frac{1}{n} \sum_{i=1}^n |W_i| h\{\text{sgn}(W_i) \times f(\mathbf{X}_i)\} + \frac{1}{2} \lambda_n \|f\|^2, \quad (9)$$

where λ_n is a tuning parameter and $\|f\|$ is some norm for f . To solve this minimization problem, we represent the ramp loss function h as the difference between two convex functions as shown in Figure 1(b): $h(u) = h_1(u) - h_2(u)$, where $h_1(u) = (1/2 - u)_+$ and $h_2(u) = (-1/2 - u)_+$. This allows application of a difference of convex functions algorithm (DCA) for minimization as will be presented in the following section. Next we consider derivation of treatment-selection rules based on linear and nonlinear marker combinations consecutively.

2.2 Identification of Linear Marker Combinations

First, we consider optimizing a linear marker combination $f(\mathbf{X}) = b + \mathbf{X}^T \boldsymbol{\beta}$. We propose to solve for b and $\boldsymbol{\beta}$ by minimizing the weighted sum of ramp loss (9) with L_2 -norm penalty

$$\frac{1}{n} \sum_{i=1}^n |w_i| h\{z_i (b + \mathbf{x}_i^T \boldsymbol{\beta}) < 0\} + \frac{1}{2} \lambda_n \|\boldsymbol{\beta}\|_2^2, \quad (10)$$

with $z_i = \text{sgn}(w_i)$.

Let $\eta_i = z_i f(\mathbf{x}_i) = z_i (b + \mathbf{x}_i^T \boldsymbol{\beta})$, and write $h(u) = h_1(u) - h_2(u)$, minimization of (10) can be carried out in the following steps:

Step 1. Start with an initial guess for $\boldsymbol{\eta}$ and call it $\boldsymbol{\eta}^0$.

Step 2. Solve

$$\hat{\boldsymbol{\beta}} = \underset{\boldsymbol{\beta}}{\text{argmin}} \sum_{i=1}^n |w_i| \left\{ h_1(\eta_i) - \hat{h}_2(\eta_i, \eta_i^0) \right\} / n + \lambda_n \|\boldsymbol{\beta}\|_2^2 / 2,$$

where $\hat{h}_2(\eta_i, \eta_i^0) = h_2(\eta_i^0) + h_2'(\eta_i^0) \eta_i$, and $h_2'(u) = -I(u < -1/2)$ is the first derivative of $h_2(u)$ with respect to u , except 0 at $u = -1/2$.

Step 3. Compute $\boldsymbol{\eta}^0$ and go back to step 2 until the change in the penalized weighted sum of ramp loss is less than a pre-specified threshold.

In Step 2 of the algorithm above, we solve a convex optimization problem. As h_1 and $\hat{h}_2$ are not smooth functions, a standard optimization strategy is to convert it to a constrained smooth optimization problem by introducing slack variables ξ_i to replace h_1 and adding two sets of constraints. This leads to

$$\min_{\xi, \beta, b} \frac{1}{n} \sum_{i=1}^n |w_i| \{ \xi_i - h'_2(\eta_i^0) \eta_i \} + \frac{1}{2} \lambda_n \|\beta\|_2^2 \quad (11)$$

subject to $\xi_i \geq \frac{1}{2} - \eta_i$, and $\xi_i \geq 0$,

where we have dropped terms void of ξ or β . The constraints in the above optimization problem are challenging to work with; therefore we seek to find its dual problem via the Lagrange method. The primary Lagrangian is

$$L_p = \frac{1}{n} \sum_{i=1}^n |w_i| \{ \xi_i - h'_2(\eta_i^0) \eta_i \} + \frac{1}{2} \lambda_n \|\beta\|_2^2 - \frac{1}{n} \sum_{i=1}^n \alpha_i \left(\xi_i + \eta_i - \frac{1}{2} \right) - \frac{1}{n} \sum_{i=1}^n \gamma_i \xi_i,$$

where α and γ are vectors of non-negative Lagrange multipliers, corresponding to the two sets of constraints in (11). As detailed in Web Supplementary Appendix C, this leads to the dual problem of

$$\min_{\alpha} \alpha^T Q \alpha - \{ \mathbf{1} - 2Q h'_{2,w}(\eta^0) \}^T \alpha$$

subject to $0 \leq \alpha_i \leq |w_i|$, and $\sum_{i=1}^n \alpha_i z_i + \sum_{i=1}^n z_i |w_i| h'_2(\eta_i^0) = 0$, (12)

where Q is a square matrix whose $[i, i']^{th}$ element is $\langle z_i \mathbf{x}_i, z_{i'} \mathbf{x}_{i'} \rangle / (n \lambda_n)$, $\mathbf{1} = \{1\}_{i=1}^n$, and $h'_{2,w}(\eta^0) = \{|w_i| h'_2(\eta_i^0)\}_{i=1}^n$. The optimization problem (12) has a set of simple box constraints and can be solved via many quadratic programming methods. As shown in Web Supplementary Appendix C, we have $\beta = \sum_{i=1}^n z_i \{ |w_i| h'_2(\eta_i^0) + \alpha_i \} \mathbf{x}_i / (n \lambda_n)$. The estimate of b can be found by averaging $\{ 1/2 z_i - \mathbf{x}_i^T \beta = 1/2 z_i - Q \times \{ h'_{2,w}(\eta) + \alpha \} / z_i : |w_i| > \alpha_i > 0 \}$. Having found the best combination, the treatment-selection rule for a new observation $\mathbf{x}_{new}$ will be $A(\mathbf{x}_{new}) = I\{f(\mathbf{x}_{new}) > 0\}$ where $f(\mathbf{x}_{new}) = b + \mathbf{x}_{new}^T \beta$.

2.3 Identification of Nonlinear Marker Combinations

As the most optimal marker combination for treatment selection may not be among linear combinations of the input markers, we may wish to enlarge the feature space via basis expansion and find the best linear combination in the enlarged space. One such approach is via the so called “kernel methods”. Let $\phi_i = \phi(\mathbf{x}_i)$ denote the feature vector for subject i in the enlarged feature space. For example, with two markers $\{x_1, x_2\}$ to derive a treatment-selection rule, we may wish to go beyond their linear combinations and look at the second order trends as well as the interaction between the two markers. Then ϕ_i in the enlarged feature space, will equal $\{x_{1i}, x_{2i}, x_{1i}^2, x_{2i}^2, x_{1i}x_{2i}\}$. We can specify a kernel K corresponding to the inner product in the mapping ϕ , i.e., $K(\mathbf{x}_i, \mathbf{x}_j) = \langle \phi_i, \phi_j \rangle$. Conversely, for any continuous, symmetric, and positive semidefinite kernel function $K: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$, Mercer theorem (Vapnik, 1995) shows that there exists such a mapping ϕ such that $K(\mathbf{x}_i, \mathbf{x}_j) = \langle \phi_i, \phi_j \rangle$. Examples of common nonlinear kernels include the polynomial kernel with d^{th} degree $K(\mathbf{x}_i, \mathbf{x}_j) = (1 + \langle \mathbf{x}_i, \mathbf{x}_j \rangle)^d$, the radial basis function (RBF) kernel $K(\mathbf{x}_i, \mathbf{x}_j) = \exp(-\gamma \|\mathbf{x}_i -$

$\mathbf{x}_j\|^2)$ with γ a tuning parameter, and the identical-by-state (IBS) kernel

$$K(\mathbf{x}_i, \mathbf{x}_j) = \sum_{p=1}^P \{2 \times I(x_{ip} = x_{jp}) + I(|x_{ip} - x_{jp}| = 1)\} / (2p) \text{ for } x_{ip} = 0, 1, 2 \text{ (Wessel and Schork, 2006).}$$

Let $\mathcal{H}_K$ be the associated reproducing kernel Hilbert space (RKHS). We generalize equation (10) to find a combination of markers $f \in \mathcal{H}_K$ through the minimization of the penalized

weighted sum of ramp loss $\sum_{i=1}^n |w_i| h\{z_i f(\mathbf{x}_i)\} / n + \lambda_n \|f\|^2 / 2$, with $\|\cdot\|$ the norm in $\mathcal{H}_K$.

This is the same as (10), when K is the linear kernel. The algorithm developed in Section 2.2 can be adapted to solve the minimization problem. In step 1 we can let $\eta_i \leftarrow 0$. In step 2, the dual problem that needs to be solved has the same form as (12), where $\mathbf{Q}$ is now a square

matrix whose $[i, i']^{th}$ element is $\langle z_i \phi_i, z_{i'} \phi_{i'} \rangle / (n \lambda_n) = z_i K(\mathbf{x}_i, \mathbf{x}_{i'}) z_{i'} / (n \lambda_n)$. Having solved the dual problem, it is not necessary to translate it back to the primal problem because step 3

calls for $\boldsymbol{\eta}$. With $\boldsymbol{\beta} = \sum_{i=1}^n z_i \{ |w_i| h'_2(\eta_i^0) + \alpha_i \} \phi_i / (n \lambda_n)$, we have

$\boldsymbol{\eta} = \mathbf{z} \cdot (\mathbf{b} + \boldsymbol{\Phi} \hat{\boldsymbol{\beta}}) = \mathbf{z} \cdot \hat{\mathbf{b}} + \mathbf{Q} \times \{ \mathbf{h}'_{2,w}(\boldsymbol{\eta}^0) + \boldsymbol{\alpha} \}$, where $\cdot$ is the element-wise multiplication operator, and $\boldsymbol{\Phi} = (\boldsymbol{\phi}_1, \dots, \boldsymbol{\phi}_n)^T$. This shows that even though $\boldsymbol{\phi}_i$ may be infinite-dimensional, η_i always has a finite-dimensional representation. As shown in Web Supplementary Appendix C, the marker combination score for a new data $\mathbf{x}_{new}$ is

$$f(\mathbf{x}_{new}) = \mathbf{b} + (\boldsymbol{\phi}_{new})^T \boldsymbol{\beta} = \mathbf{b} + \sum_{i=1}^n \{ \mathbf{h}'_{2,w}(\boldsymbol{\eta}) + \alpha_i \} z_i K(\mathbf{x}_{new}, \mathbf{x}_i) / (n \lambda_n).$$

In practice, having the flexibility of modeling nonlinear marker effect is important for achieving good treatment-selection performance. This is especially true in the presence of nonzero treatment burden. Using the GLM model as an example: suppose the disease risk conditional on the markers and treatment follows $g\{P(Y = 1|\mathbf{X}, T)\} = \beta_0 + \beta_1 T + \beta_2 \mu(\mathbf{X}) + \beta_3 \nu(\mathbf{X}) T$ with $\mu(\mathbf{X})$ and $\nu(\mathbf{X})$ being some functions of $\mathbf{X}$. Based on (3), the optimal treatment-selection rule is to treat if the difference between $Risk_0(\mathbf{X}) = P(Y = 1|\mathbf{X}, T = 0)$ and $Risk_1(\mathbf{X}) = P(Y = 1|\mathbf{X}, T = 1)$ is greater than the treatment/disease burden ratio δ . Using the GLM model the optimal treatment-selection rule, given $\delta = 0$, is to treat if $\beta_1 + \beta_3 \nu(\mathbf{X}) > 0$ and to not treat otherwise, since $\text{sgn}\{Risk_0(\mathbf{X}) - Risk_1(\mathbf{X})\} = \text{sgn}\{\beta_1 + \beta_3 \nu(\mathbf{X})\}$. In other words, only the main effect of treatment and the marker-treatment interaction need to be modeled correctly to derive the optimal $A(\mathbf{X})$ for $\delta = 0$. This is no longer true, however, for positive δ , since the sign of $Risk_0(\mathbf{X}) - Risk_1(\mathbf{X}) - \delta$ also depends on the intercept term and the main effect of the marker in the risk model for nonlinear link function g (which makes it more challenging to closely model the marker combination). Note that if there exists an underlying monotone transformation q such that $q\{Risk_0(\mathbf{X}) - Risk_1(\mathbf{X})\} = c + \gamma(\mathbf{X})$ for some constant c , then the optimal rule is to recommend treatment if $c + \gamma(\mathbf{X}) > q(\delta)$ or equivalently if $c - q(\delta) + \gamma(\mathbf{X}) > 0$. This optimal rule belongs to the class of treatment-selection rules we study in this paper.

2.4 Theoretical Results

We consider the problem of minimizing $E[|W| \times I\{\text{sgn}(W) \times f(\mathbf{X}) \neq 0\}]$, the expectation of (4) with respect to a measurable function f , where W is a random variable indicating individual-specific weight. We make the following assumptions: (i) f is a measurable

function: $\mathcal{X} \rightarrow \mathbb{R}$; (ii) $Y \in \mathbb{R}$ is bounded; (iii) $|W|$ for $W \in \mathbb{R}$ is bounded above. Fisher consistency of the treatment-selection rule based on minimizing the weighted sum of the ramp loss can be derived as shown in Theorem 1.

Theorem 1—For any measurable function f , if $\hat{f}$ minimizes $E[|W| \times h\{\text{sgn}(W) \times f(\mathbf{X})\}]$, then $\hat{f}$ also minimizes $E[|W| \times I\{\text{sgn}(W) \times f(\mathbf{X}) - 0\}]$.

Theorem 1 states that the treatment-selection rule minimizing the weighted sum of ramp loss will also minimize the weighted sum of 0–1 loss. The approximation of the 0–1 loss with the ramp loss does not affect the best treatment-selection performance. Proof of Theorem 1 is shown in Web Supplementary Appendix D.

Furthermore, let $\hat{f}_n$ be the penalized ramp estimator of f proposed in Section 2.1 based on a sample size of n ; one can show that the weighted sum of ramp loss estimated based on $\hat{f}_n$ is consistent for that based on f as presented in Theorem 2.

Theorem 2—For a sequence of $\lambda_n > 0$ such that $\lambda_n \rightarrow 0$ and $\lambda_n n \rightarrow \infty$, we have that in probability $\lim_{n \rightarrow \infty} \{\theta_h(\hat{f}_n) - \inf_{f \in H_k} \theta_h(f)\} = 0$, where $\bar{\mathcal{H}}_k$ denotes the closure of $\mathcal{H}_k$ and $\theta_h(f) = E[|W| \times h\{\text{sgn}(W)f\}]$.

Proof of Theorem 2 is presented in Web Supplementary Appendix E.

2.5 Alternative Estimator of Marker Combinations

In this section, we briefly discuss an alternative way to minimizing the weighted sum of 0–1 loss in (4) based on the weighted support vector machine technology. Since

$$\sum_{i=1}^n W_i \times I\{f(\mathbf{X}_i) \leq 0\} \propto \sum_{i=1}^n |W_i| \times I[\text{sgn}(W_i) \neq \text{sgn}\{f(\mathbf{X}_i)\}],$$

minimizing the weighted sum of 0–1 loss $\sum_{i=1}^n W_i \times I\{f(\mathbf{X}_i) \leq 0\}$ can be translated into an equivalent problem of minimizing a weighted sum of classification error

$\sum_{i=1}^n W_i^* \times I[\text{sgn}(W_i) \neq \text{sgn}\{f(\mathbf{X}_i)\}]$, where W_i^* , the weight for the new classification problem, equals $|W_i|$. In other words, our minimization problem can be formulated into a problem of minimizing the weighted sum of classification error where

$Y_i^* = \text{sgn}(W_i) = \{-1, 1\}$ is the true binary classes, $\text{sgn}\{f(\mathbf{X}_i)\}$ is the predicted binary class based on $\mathbf{X}$, and W_i^* is the subject-specific weight. This type of weighted classification error problem can thus be resolved using the weighted support vector machine (Lin and Wang, 2002) based on $\mathbf{X}_i$, Y_i^* , and W_i^* (sketched in Web Supplementary Appendix F). The outcome-dependent estimator of Zhao et al. (2012) is a special case of this class with weights W^* corresponding to the case-only weights W_1 (5) in the scenario where $\delta = 0$.

This alternative approach is essentially approximating the 0–1 loss with a hinge loss function $g(u) = \max(0, 1 - u)$, and minimizing the weighted sum of hinge loss. As shown in Figure 1(a), the hinge loss does not approximate the 0–1 loss adequately. Use of a suboptimal approximation could compromise the performance of the estimated marker combinations. Next, in simulation studies, we include the weighted support vector machine

estimator as a comparison estimator. We also assess the impact of different choices of weights relative to the outcome-dependent weight adopted in Zhao et al. (2012).

3. Simulation

In this section, we investigate the performance of the proposed estimator that minimizes the penalized weighted sum of ramp loss for treatment selection (W-RAMP). For three different simulation settings of 1:1 randomized trial, we study two common kernels for combining two biomarkers X_1 and X_2 : the linear kernel for finding linear marker combinations and the RBF kernel for identifying nonlinear combinations. For each simulation setting, we consider the treatment/disease burden ratio $\delta = 0$ and 0.01, and evaluate four different sets of weights: a) the case-only weights (W_1); b) the control-only weights (W_2); c) the average case+control weights (W_3); and d) the double-robustness weights (W_4). Performance of the W-RAMP estimators is compared with the common linear logistic regression model and the weighted support vector machine (W-SVM) estimators described in Section 2.5. In the linear logistic regression, the risk of Y is modeled as a function of the main effect for treatment status and for each biomarker, as well as the interaction between treatment status and each biomarker. This is also the risk model used for constructing the double-robustness weights. For comparison, we also applied the genetic algorithm (W-GENOUD) to find linear marker combinations that minimize the weighted sum of 0–1 loss (4) as did in Zhang et al. (2012a).

In each Monte Carlo simulation, a training sample of size $n = 100$ is generated; 5-fold cross-validation is used to identify the optimal tuning parameter λ_n and γ (for RBF kernel) within a range of 0 to 5 after $\sum_{i=1}^n |W_i|/n$ is scaled to be 1; a treatment-selection rule $\hat{A}(\mathbf{X})$ is then estimated from the training sample using the selected tuning parameter. To evaluate the performance of the estimated rule, a test set of $n = 100,000$ is generated in each simulation, based on which we estimate θ with

$$\sum_{i=1}^n (1-T_i) \times Y_i \times \{1 - \hat{A}(\mathbf{X}_i)\} / n_0 + \sum_{i=1}^n T_i \times Y_i \times \hat{A}(\mathbf{X}_i) / n_1 + \delta \times \sum_{i=1}^n \hat{A}(\mathbf{X}_i) / n.$$

Distribution of the test set performance over 500 Monte-Carlo simulations is evaluated. For ease of comparison, here we present the percent reduction in θ using a marker-based rule compared to the optimal treatment strategy in the absence of marker information. The latter is defined as treating all if the difference in disease prevalence between the untreated and the treated is larger than δ , or treating none otherwise.

Setting 1 includes two sub-settings. In the first sub-setting (1A), we simulate two biomarkers (X_1, X_2) from a bivariate normal distribution with mean 0, variance 1 each, and correlation 0.2. Risk of Y conditional on X_1, X_2 and T follows a linear logistic model $\text{logit}P(Y=1|X_1, X_2, T) = -1.6 - X_1 - X_2 - T \times (2.5 + 2.5X_1 + 2.5X_2)$. In the second sub-setting (1B), 10% outliers are added to the data, where the outliers consist of independent markers X_1, X_2 each with mean 2 and variance 2, with the risk of Y conditional on X_1, X_2 and T following a logistic model

$\text{logit}P(Y=1|X_1, X_2, T) = 2 + 1.1X_1^3 + 1.1X_2^3 - T \times (1.3X_1^2 + 1.3X_2^2 + 1.3X_1X_2)$. Disease prevalences among untreated and treated are 0.244 and 0.236 respectively for setting 1A and 0.318 and 0.307 respectively for setting 1B.

Setting 2 also includes two sub-settings. In the first sub-setting (2A), we simulate two biomarkers (X_1, X_2) from a bivariate normal distribution with mean 0, variance 1 each, and correlation 0.2. Risk of Y conditional on X_1, X_2 and T follows a logistic model

$\text{logit}P(Y=1|X_1, X_2, T) = -1.5 + 0.2X_1 - 0.2X_2 + T \times (-3 - X_1 - X_2 + X_1^2 + X_2^2)$. In the second sub-setting (2B), 10% outliers are added where the outliers consist of independent markers X_1, X_2 each with mean 2 and variance 1.5, with the risk of Y conditional on X_1, X_2 and T following a logistic model

$\text{logit}P(Y=1|X_1, X_2, T) = -1.5 + 0.2X_1 - 0.2X_2 + T \times (-3 - X_1 - X_2 - 0.1X_1^2 - 0.1X_2^2 + 0.4X_1^3 + 0.4X_2^3)$. Disease prevalences among untreated and treated are 0.185 and 0.175 respectively for setting 2A and 0.192 and 0.190 respectively for setting 2B.

In setting 3, we simulate two independent biomarkers (X_1, X_2) each from a mixture of 90% uniform(-1,1) and 10% uniform(0,2). Risk of Y conditional on X_1, X_2 and T follows $\text{logit}P(Y=1|X_1, X_2, T) = -1 + T \times [0.5 - 6X_1X_2 \times I\{(X_1 > 0 \& X_2 > 0)\}(X_1 < 0 \& X_2 < 0)\}]$. Disease prevalences among untreated and treated are 0.269 and 0.256 respectively.

Tables 1 shows simulation results for $\delta = 0$. In setting 1A, where the linear logistic model holds, the marker combination based on the linear logistic model achieves the largest treatment-selection benefit with the highest precision. The performance of the linear W-RAMP estimator is inferior to the linear logistic model, but significantly better than the linear W-SVM estimator (p-value = 0.001 comparing θ based on the paired t-test for each set of weights). Moreover, the performance of the linear W-SVM estimator is sensitive to the choice of weights, with the double-robustness weights greatly outperforming the other three sets of weights including the case-only weights. In contrast, the performance of the linear W-RAMP estimator is fairly insensitive to the choice of weights. When the RBF kernel is used, the double-robustness and the case-only weights work best for both the W-RAMP and the W-SVM estimators; the W-RAMP estimator significantly outperforms the W-SVM estimator with the double-robustness weights (p-value=0.008 comparing θ based on the paired t-test) while the two have comparable performance with respect to the other three weights.

With outliers added (settings 1B), linear marker combinations identified by the W-RAMP estimator have much improved performance compared to either the linear logistic model or the W-SVM. Again, the performance of the W-RAMP is insensitive to the choice of weights. Nonlinear marker combinations identified using the RBF kernel have improved performance compared to linear marker combinations for W-SVM, with the performance comparable between the W-RAMP and the W-SVM.

In setting 2A, where interactions exist between treatment and polynomial marker effect, the linear W-RAMP estimator has a slightly better treatment-selection performance compared to the linear logistic model, whereas the linear W-SVM estimator has the worst performance. With the RBF kernel, both the W-RAMP and the W-SVM estimators have slightly improved performance compared to linear marker combinations, especially for the case-only weights and the double-robustness weights. The performance is comparable between the W-RAMP and the W-SVM. When 10% outliers are added to the data (setting 2B), among linear marker combinations, the W-RAMP estimator significantly outperforms the linear logistic model or

the linear W-SVM. Using the RBF kernel allowing for nonlinear effect leads to a large improvement in treatment-selection benefit compared to the linear kernel with the performance comparable between the W-RAMP and the W-SVM.

In setting 3, where the optimal marker combination for treatment selection is highly nonlinear, the linear logistic model and the linear W-SVM lead to marker combinations that perform worse than the strategy of treating all. The linear W-RAMP estimator, in contrast, still has better performance compared to treating all. The RBF kernel improves the performance over linear kernel for both the W-RAMP and the W-SVM estimators with performance similar between the two.

For finding linear combinations, the relative performance of the W-RAMP and the W-GENOUD varies with the weights and settings, but overall the two seem to have comparable performance. Computation times for W-GENOUD, W-RAMP, and W-SVM are around 4.6, 0.3, and 0.1 sec respectively for a training size of 100.

Tables 2 shows simulation results for $\delta = 0.01$. The pattern for comparing different estimators is fairly similar to that in Table 1. In general, linear marker combinations derived from the proposed W-RAMP estimator have greater performance compared to those from the linear W-SVM. With the presence of nonlinear effect or outlying observations, using a nonlinear kernel such as the RBF kernel improves the performance of the W-RAMP or the W-SVM. The performance of the two algorithms using the RBF kernel are comparable in various simulations when tuning is performed using 5-fold CV.

In Supplementary Tables 1 and 2 we show additional simulation results where an independent tuning set of size 1000 is generated in each Monte Carlo simulation, reflecting an “ideal” setting where there is little uncertainty in choosing tuning parameters. Under this kind of “ideal” tuning, appreciable improvement in performance of both the W-RAMP and the W-SVM are observed compared to 5-fold CV for tuning, with the W-RAMP in general performing significantly better than the W-SVM using both the linear and the RBF kernel. The potential of using better tuning methods to improve the performance of the W-RAMP merits further investigation.

4. Data Example

We illustrate the proposed estimators of marker combinations with a real example from an HIV vaccine study. As a prevention measure, vaccines typically have low cost related to treatment, thus we set $\delta = 0$. The example comes from the RV144 Thailand HIV vaccine trial, the first HIV vaccine trial that demonstrates a positive vaccine efficacy in preventing HIV infection (with an estimated hazard ratio of vaccine vs placebo of 31.2% and p-value 0.04). The trial included 16,402 participants aged 18–30 who were 1:1 randomized into a vaccine and a placebo arm (Rerks-Ngarm et al., 2009).

A followup RV144 host-genetic study was conducted to investigate the effect of genotypes of Fc receptor gene on vaccine efficacy. Overall 148 single nucleotide polymorphisms (SNPs) covering five Fc- γ receptor genes Fc- γ R2a, Fc- γ R2b, Fc- γ R2c, Fc- γ R3a, and Fc- γ R3b, and 42 SNPs covering the Fc- α receptor gene were genotyped on 125 cases (74

placebo recipients and 51 vaccine recipients), and 225 controls (20 placebo recipients and 205 vaccine recipients). After exclusion of SNPs with minor allele frequency less than 5% and SNPs highly correlated with one another (with Pearson correlation greater than 0.80), four SNPs, each categorized into a binary variable, were identified to have significant interactions with the vaccination status based on the univariate case-only analysis (Li et al., 2014).

Here we explore combinations of the four SNPs to recommend the use of vaccine in order to minimize the rate of HIV infection in the population. We consider the following five treatment strategies: (i) the strategy of not vaccinating anyone; (ii) the strategy of vaccinating everyone; (iii) the strategy of selective vaccinating based on a linear logistic model selected using the Akaike Information Criteria (AIC) defined as $2 \times \text{number of parameters} - 2 \times \log(\text{likelihood})$, where the predictors considered are the main effects for treatment and for each SNP and the interaction between treatment and each SNP; and iv) the W-RAMP estimator for four different set of weights (where the linear logistic risk model as in (iii) is used to construct the double-robustness weights); and (v) the W-SVM estimator. For the W-RAMP and the W-SVM, we consider both the linear and the IBS kernel, and use 5-fold CV for tuning parameter selection. Based on the entire dataset, the tuning parameter identified, and corresponding estimated treatment-selection rules using the linear W-RAMP and the linear W-SVM are presented (Web supplementary Table 3). Also presented is the estimated treatment-selection rule based on the logistic regression model.

In Table 3, we show estimated performance of these different strategies for deriving treatment-selection rules. Naïve estimates of θ are obtained by applying the treatment-selection rule estimated using the full dataset to the same data. Moreover, a random cross-validation procedure is performed for the computation of θ correcting for overfitting bias. In particular, the data are randomly split into five folds stratified on the disease and the treatment status, 4 folds for training and the remaining 1 fold for testing. Based on each random training subset, a 5-fold CV is performed to select the tuning parameter specific for the subset, which is then used to derive a treatment-selection rule using the random training subset; the estimated rule is then applied to the random test subset to estimate θ . The procedure is repeated 100 times with the average θ computed. In Table 3, we present both the naïve and the cross-validated estimates of θ . Percentile bootstrap confidence intervals are provided for those measures. In the bootstrap procedure, 100 resamplings are performed stratified on the disease and the treatment status; the procedure of computing the naïve and the cross-validated θ estimates is then applied to each resampled data.

From Table 3, the strategy of treating none and treating all lead to an estimated HIV infection rate of 9.26 and 6.41 per 1000 persons respectively, consistent with the positive vaccine efficacy we observed in the RV144 trial. Severe over-optimism is observed in naïve estimates of marker-based treatment-selection performance. For most methods, the naïve estimates show reduced HIV infection rate compared to the treating all strategy. After using cross-validation procedure to correct for the overfitting bias, the proposed W-RAMP estimator with either the linear or the IBS kernel and the case-only weights lead to reduced HIV infection rates: 5.74 per 1000 persons (95% CI: 4.02 to 6.71) for linear and 5.68 per 1000 persons (95% CI: 4.11 to 6.75) for IBS, compared to 6.57 per 1000 persons (95% CI:

4.94 to 7.82) for the treating all strategy. For comparison, the logistic model selected by AIC has θ estimate 6.80 (95% CI: 5.41 to 7.99) per 1000 persons, which is 3.5% above treating-all; the best W-SVM estimators (with the linear or IBS kernel and the case-only weights) have θ values fairly close to treating all. Using the W-RAMP with the case-only weights, 58.2% (95% CI: 45.8% to 73.8%) and 58.6% (95% CI: 44.2% to 73.3%) subjects will be recommended for vaccination based on cross-validation, using the linear and the IBS kernel respectively.

5. Discussion

We proposed a general framework for deriving the optimal treatment-selection strategy using biomarkers, towards the goal of minimizing the total disease and treatment burden to the population. The proposed estimation method has an attribute similar to the work of Zhang et al. (2012a) and Zhao et al. (2012) in that the optimal marker combination is identified through directly minimizing an unbiased estimate of the criterion function, and thus is more robust to model misspecification than standard regression-based methods.

Direct minimization of an estimated criterion function poses high computational challenges due to the involvement of a non-convex 0–1 loss function. To address this issue, we proposed the use of the difference of convex functions algorithm to minimize a penalized weighted sum of ramp loss. The ramp loss serves as an approximation to the 0–1 loss. It is demonstrated to 1) be a more precise approximation than the hinge-loss approximation underlying the standard support vector machine methods; and 2) lead to better or comparable estimators of marker combinations in various scenarios. For binary disease outcomes, we found the performance of the ramp estimator to be relatively insensitive to the choice of weights, with the case-only and the double-robustness weights performing best overall, and we therefore recommend the use of the two kinds of weights in practice. The good performance of the case-only weights is appealing when the criterion function to be optimized is the average disease rate, since the computation of the case-only weights in this scenario requires only measurements of markers among case samples and is thus cost-effective under limited resources.

As mentioned in Section 2.1, unlike the case of $\delta = 0$, in the presence of positive δ , the optimal treatment-selection rule can not be represented as a simple function of the main treatment effect and the marker-treatment interaction in the disease risk model. This makes the modeling of the marker combination more challenging and calls for greater flexibility in estimation methods. Our proposed kernel-estimator is appealing given its ability to allow for flexible modeling of marker combinations. In practice, it is necessary to put the disease burden and the treatment burden on the same scale to determine δ . For example, as in Rapsomaniki et al. (2012), if one can associate a monetary cost with the reduction in the rate of the targeted disease through treatment, and a monetary cost with treatment (including the cost of conducting the procedure and the cost due to secondary events), then δ can be computed as the rate of the latter relative to the former. In practice, a series of δ might also be chosen for sensitivity analysis. More detailed investigation is warranted for evaluating the performance of our estimators in the presence of positive cost ratios.

Finally, note that Janes et al. (2013) considers a more general decision-theoretical approach that allows the burden of disease and treatment to vary with biomarkers. The framework in Vickers et al. (2007) is a special case of this general formulation assuming these burdens do not depend on the marker value. The method developed in the paper, however, can be readily extended to address the general situation through specification of utility parameters that character the association between biomarkers and the disease and treatment burdens.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

The authors thank Dr. Dan Geraghty from the Fred Hutchinson Cancer Center for generating the genetics data, and Drs. Sue Li and Peter Gilbert for pre-processing and preliminary analysis of the SNP data. We thank the editor, the AE, and the referees for their help in improving this paper. This work was supported by NIH grants R01 GM106177-01, 2R37AI05465-10, U01 CA086368, and P01 CA053996.

References

- Baker SG, Kramer BS, Sargent DJ, Bonetti M. Biomarkers, subgroup evaluation, and clinical trial design. *Discovery Medicine*. 2012; 13(70):187–192. [PubMed: 22463794]
- Foster JC, Taylor JMG, Ruberg SJ. Subgroup identification from randomized clinical trial data. *Statistics in Medicine*. 2011; 30(24):2867–2880. [PubMed: 21815180]
- Janes H, Pepe M, Huang Y. A general framework for evaluating markers used to select patient treatment. *Medical Decision Making*. 2013;10.1177/0272989X13493147
- Li SBGP, Tomaras GD, Kijak G, Ferrari G, Thomas R, et al. Association of fcgr2c polymorphisms with vaccine efficacy and correlates of hiv-1 infection risk in the rv144 trials. *Journal of Clinical Investigation*. 2014 under revision.
- Lin C, Wang S. Fuzzy support vector machines. *Neural Networks, IEEE Transactions on*. 2002; 13(2): 464–471.
- Liu Y, Shen X, Doss H. Multicategory ψ -learning and support vector machine: computational tools. *Journal of Computational and Graphical Statistics*. 2005; 14(1)
- Lu W, Zhang HH, Zeng D. Variable selection for optimal treatment decision. *Statistical Methods in Medical Research*. 2013; 22(5):493–504. [PubMed: 22116341]
- McKeague, IW.; Qian, M. *Statistica Sinica*. 2013. Estimation of treatment policies based on functional predictors. In Press
- Orellana L, Rotnitzky A, Robins JM. Dynamic regime marginal structural mean models for estimation of optimal dynamic treatment regimes, part i: main content. *The International Journal of Biostatistics*. 2010; 6(2)
- Qian M, Murphy S. Performance guarantees for individualized treatment rules. *The Annals of Statistics*. 2011; 39(2):1180–1210.
- Rapsomaniki E, White I, Wood A, Thompson S. A framework for quantifying net benefits of alternative prognostic models. *Statistics in Medicine*. 2012; 31(2):114–130. [PubMed: 21905066]
- Rerks-Ngarm S, Pitisuttithum P, Nitayaphan S, Kaewkungwal J, Chiu J, Paris R, et al. Vaccination with ALVAC and AIDSVAX to prevent HIV-1 infection in Thailand. *New England Journal of Medicine*. 2009; 361(23):2209–2220. [PubMed: 19843557]
- Rubin DB. Comment on “randomization analysis of experimental data: The fisher randomization test” by d. basu. *Journal of the American Statistical Association*. 1980; 75(371):591–593.
- Song X, Pepe MS. Evaluating markers for selecting a patient’s treatment. *Biometrics*. 2004; 60(4): 874–883. [PubMed: 15606407]
- Vapnik, V. *The Nature of Statistical Learning Theory*. Springer Verlag; 1995.

- Vickers A, Kattan M, Sargent D. Method for evaluating prediction models that apply the results of randomized trials to individual patients. *Trials*. 2007; 8(1):14. [PubMed: 17550609]
- Wessel J, Schork NJ. Generalized genomic distance-based regression methodology for multilocus association analysis. *The American Journal of Human Genetics*. 2006; 79(5):792–806.
- Wu Y, Liu Y. Robust truncated hinge loss support vector machines. *Journal of the American Statistical Association*. 2007; 102(479):974–983.
- Zhang B, Tsiatis A, Laber E, Davidian M. A robust method for estimating optimal treatment regimes. *Biometrics*. 2012a; 68(4):1010–1018. [PubMed: 22550953]
- Zhang B, Tsiatis AA, Davidian M, Zhang M, Laber EB. Estimating optimal treatment regimes from a classification perspective. *Stat*. 2012b; 1(1):103–114. [PubMed: 23645940]
- Zhao Y, Zeng D, Rush AJ, Kosorok MR. Estimating individualized treatment rules using outcome weighted learning. *Journal of the American Statistical Association*. 2012; 107(499):1106–1118. [PubMed: 23630406]

6. Web Supplementary Appendix

Web Appendices and Tables referenced in Sections 2, 2.1–2.5, 3, and 4, and example R code for implementing the proposed method are available with this paper at the Biometrics website on Wiley Online Library.

Figure 1.

(a) Different approximations of the 0-1 loss function. (b) Representing the ramp loss as the difference of two convex functions.

Table 1

Percent reduction in θ using marker-based rule compared to the best treatment strategy in the absence of marker information, for $\delta = 0$, training size $n = 100$. Five-fold CV is used for selecting tuning parameters.

Scenario	ρ_0	ρ_1	Method	Weights for Ramp and SVM							
				Case-only		Control-only		CC ^a		DR ^b	
				Mean	SE	Mean	SE	Mean	SE	Mean	SE
I-A	0.244	0.236	Linear logistic	21.15	(0.25)						
			W-GENOUD, Linear	10.81	(0.34)	3.89	(0.34)	8.31	(0.34)	18.52	(0.34)
			W-Ramp, Linear	14.12	(0.43)	13.90	(0.42)	14.94	(0.42)	14.73	(0.44)
			W-SVM, Linear	7.09	(0.44)	0.98	(0.25)	2.40	(0.29)	12.08	(0.48)
			W-Ramp, RBF	8.80	(0.38)	5.59	(0.33)	6.72	(0.31)	11.01	(0.39)
			W-SVM, RBF	8.78	(0.39)	6.06	(0.34)	6.66	(0.32)	9.82	(0.37)
I-B	0.318	0.307	Linear logistic	1.56	(0.35)						
			W-GENOUD, Linear	4.71	(0.37)	2.15	(0.37)	4.38	(0.37)	4.95	(0.37)
			W-Ramp, Linear	5.23	(0.35)	4.59	(0.36)	5.11	(0.38)	5.19	(0.36)
			W-SVM, Linear	-0.58	(0.16)	-0.09	(0.18)	-0.59	(0.16)	-0.43	(0.17)
			W-Ramp, RBF	4.53	(0.27)	3.12	(0.24)	4.35	(0.24)	5.04	(0.26)
			W-SVM, RBF	4.33	(0.26)	3.10	(0.24)	4.36	(0.24)	4.63	(0.25)
II-A	0.185	0.175	Linear logistic	30.81	(0.65)						
			W-GENOUD, Linear	31.14	(0.53)	16.03	(0.53)	26.92	(0.53)	33.22	(0.53)
			W-Ramp, Linear	31.19	(0.67)	30.41	(0.65)	32.27	(0.62)	31.44	(0.67)
			W-SVM, Linear	24.17	(0.83)	12.07	(0.74)	17.43	(0.77)	23.45	(0.83)
			W-Ramp, RBF	32.62	(0.68)	23.30	(0.7)	30.86	(0.67)	32.42	(0.69)
			W-SVM, RBF	32.47	(0.68)	22.75	(0.71)	30.45	(0.68)	32.64	(0.66)
II-B	0.192	0.19	Linear logistic	15.57	(0.68)						
			W-GENOUD, Linear	21.99	(0.53)	10.85	(0.53)	19.72	(0.53)	23.49	(0.53)
			W-Ramp, Linear	21.06	(0.57)	19.30	(0.62)	21.08	(0.59)	21.77	(0.58)
			W-SVM, Linear	12.76	(0.62)	7.55	(0.53)	9.12	(0.54)	12.51	(0.60)
			W-Ramp, RBF	31.31	(0.64)	23.94	(0.68)	29.47	(0.63)	32.62	(0.61)
			W-SVM, RBF	31.46	(0.64)	23.81	(0.67)	29.55	(0.64)	32.11	(0.60)
III	0.269	0.256	Linear logistic	-0.54	(0.23)						

Weights for Ramp and SVM											
Scenario	ρ_0	ρ_1	Method	Case-only		Control-only		CC ^a		DR ^b	
				Mean	SE	Mean	SE	Mean	SE	Mean	SE
			W-GENOUD, Linear	1.94	(0.25)	-0.06	(0.25)	1.79	(0.25)	2.31	(0.25)
			W-Ramp, Linear	1.19	(0.22)	0.10	(0.21)	0.76	(0.22)	1.16	(0.22)
			W-SVM, Linear	-0.24	(0.19)	-1.01	(0.21)	-1.20	(0.18)	-0.87	(0.18)
			W-Ramp, RBF	5.19	(0.31)	1.53	(0.25)	3.38	(0.26)	4.04	(0.28)
			W-SVM, RBF	5.17	(0.31)	1.48	(0.26)	3.39	(0.26)	4.28	(0.26)

CC^a: Case-control weights; DR^b: double-robustness weights

Percent reduction in θ using marker-based rule compared to the best treatment strategy in the absence of marker information, for $\delta = 0.01$ and training size $n = 100$. Five-fold CV is used for selecting tuning parameters.

Table 2

Scenario	ρ_0	ρ_1	Method	Weights for Ramp and SVM							
				Case-only		Control-only		CC ^a		DR ^b	
				Mean	SE	Mean	SE	Mean	SE	Mean	SE
I-A	0.244	0.246	Linear logistic	17.32	(0.37)						
			W-GENOUD, Linear	7.68	(0.37)	4.19	(0.37)	7.64	(0.37)	17.24	(0.37)
			W-Ramp, Linear	11.53	(0.41)	14.94	(0.40)	14.69	(0.40)	13.72	(0.41)
			W-SVM, Linear	6.24	(0.40)	1.41	(0.21)	3.10	(0.28)	10.96	(0.43)
			W-Ramp, RBF	5.78	(0.39)	6.88	(0.32)	6.49	(0.29)	9.10	(0.36)
I-B	0.318	0.317	W-SVM, RBF	5.62	(0.39)	7.14	(0.29)	6.08	(0.28)	8.48	(0.34)
			Linear logistic	2.60	(0.27)						
			W-GENOUD, Linear	3.92	(0.34)	2.83	(0.34)	4.86	(0.34)	5.61	(0.34)
			W-Ramp, Linear	4.60	(0.32)	5.99	(0.34)	6.09	(0.34)	5.76	(0.34)
			W-SVM, Linear	0.77	(0.15)	1.28	(0.18)	0.97	(0.15)	0.93	(0.14)
II-A	0.185	0.185	W-Ramp, RBF	3.84	(0.27)	5.10	(0.23)	5.15	(0.22)	5.53	(0.24)
			W-SVM, RBF	3.88	(0.27)	5.18	(0.21)	5.26	(0.22)	5.24	(0.22)
			Linear logistic	26.71	(0.60)						
			W-GENOUD, Linear	26.19	(0.53)	15.50	(0.53)	25.70	(0.53)	32.13	(0.53)
			W-Ramp, Linear	28.90	(0.61)	28.92	(0.63)	31.65	(0.58)	30.80	(0.60)
II-B	0.192	0.2	W-SVM, Linear	23.35	(0.76)	12.94	(0.69)	19.09	(0.72)	23.25	(0.75)
			W-Ramp, RBF	28.54	(0.69)	22.76	(0.62)	28.58	(0.61)	31.34	(0.59)
			W-SVM, RBF	29.08	(0.71)	23.14	(0.60)	28.27	(0.62)	30.57	(0.62)
			Linear logistic	12.89	(0.59)						
			W-GENOUD, Linear	15.43	(0.52)	7.09	(0.52)	15.50	(0.52)	20.37	(0.52)
III	0.269	0.266	W-Ramp, Linear	15.64	(0.58)	15.63	(0.58)	18.09	(0.56)	18.14	(0.54)
			W-SVM, Linear	9.28	(0.57)	4.35	(0.52)	6.75	(0.51)	10.14	(0.59)
			W-Ramp, RBF	24.74	(0.66)	19.66	(0.59)	25.76	(0.58)	28.04	(0.60)
			W-SVM, RBF	25.04	(0.66)	21.02	(0.57)	25.78	(0.57)	28.21	(0.58)
			Linear logistic	-0.03	(0.11)						

Weights for Ramp and SVM											
Scenario	ρ_0	ρ_1	Method	Case-only		Control-only		CC ^a		DR ^b	
				Mean	SE	Mean	SE	Mean	SE	Mean	SE
			W-GENOUD, Linear	3.97	(0.25)	2.38	(0.25)	3.96	(0.25)	4.20	(0.25)
			W-Ramp, Linear	2.41	(0.24)	1.97	(0.23)	2.32	(0.22)	2.48	(0.22)
			W-SVM, Linear	0.90	(0.17)	0.55	(0.20)	0.78	(0.17)	0.63	(0.16)
			W-Ramp, RBF	6.07	(0.37)	4.19	(0.32)	5.74	(0.27)	6.05	(0.27)
			W-SVM, RBF	5.31	(0.33)	3.26	(0.25)	5.80	(0.26)	5.97	(0.26)

CC^a: Case-control weights; DR^b: double-robustness weights

Table 3

Estimates of $\theta/1000$ and percentile bootstrap CIs for different methods of deriving SNP combinations in the RV144 host-genetics study.

Method	Weight	Naive		Cross-validated	
		Est	95% CI	Est	95% CI
Treat None		9.26	(7.45, 10.87)	8.68★	(6.99, 11.77)
Treat All		6.41	(4.94, 8.03)	6.57	(4.94, 7.82)
Linear Logistic		6.14	(4.50, 7.58)	6.80	(5.38, 7.85)
W-RAMP, Linear	W1	4.96	(3.79, 5.88)	5.74	(4.02, 6.71)
	W2	8.75	(6.40, 10.65)	8.74	(6.88, 9.91)
	W3	8.62	(6.79, 10.84)	9.08	(7.01, 10.42)
	W4	5.49	(4.70, 7.32)	6.92	(5.15, 7.58)
W-SVM, Linear	W1	4.96	(3.98, 6.41)	6.53	(4.25, 7.15)
	W2	9.41	(6.00, 9.80)	9.06	(6.80, 9.98)
	W3	9.41	(5.67, 9.80)	9.04	(6.73, 9.97)
	W4	6.14	(4.96, 6.41)	6.63	(4.99, 7.73)
W-RAMP, IBS	W1	4.96	(3.79, 6.01)	5.68	(4.11, 6.75)
	W2	8.75	(6.79, 10.72)	9.14	(7.07, 10.34)
	W3	9.41	(6.66, 11.17)	9.18	(7.04, 10.41)
	W4	6.14	(4.84, 7.72)	7.01	(5.26, 7.60)
W-SVM, IBS	W1	5.75	(4.11, 6.41)	6.54	(4.25, 7.17)
	W2	9.41	(6.00, 9.80)	9.07	(6.78, 9.82)
	W3	9.41	(6.00, 9.80)	9.04	(6.74, 9.96)
	W4	6.14	(4.96, 6.41)	6.61	(4.97, 7.73)

★ net benefit (Vickers et al., 2007) of a treatment-selection rule equals θ of treating none $(8.68/1000)$ minus θ of treating all $= (8.68 - 6.57)/1000 = 2.11/1000$; net benefit of W-RAMP, IBS $= (8.68 - 5.68)/1000 = 3/1000$. Bootstrap is performed on the efficacy trial data stratified on the treatment status for “treat none” and “treat all” and stratified on the disease and the treatment status for other algorithms.