Published in final edited form as:

Clin Trials. 2015 August ; 12(4): 348–356. doi:10.1177/1740774515580126.

Identifying Optimal Biomarker Combinations for Treatment Selection through Randomized Controlled Trials

Ying Huang^{1,2}

¹Fred Hutchinson Cancer Research Center, Seattle, WA, USA

²University of Washington, Seattle, WA, USA

Abstract

Background/Aims—Biomarkers associated with treatment-effect heterogeneity can be used to make treatment recommendations that optimize individual clinical outcomes. To accomplish this, statistical methods are needed to generate marker-based treatment-selection rules that can most effectively reduce the population burden due to disease and treatment. Compared to the standard approach of risk modeling to derive treatment-selection rules, a more robust approach is to directly minimize an unbiased estimate of total disease and treatment burden among a pre-specified class of rules. This problem is one of minimizing a weighted sum of 0-1 loss function, which is computationally challenging to solve due to the non-smoothness of 0-1 loss. Huang and Fong,¹ among others, proposed a method that uses the Ramp loss to approximate the 0-1 loss and solves the minimization problem through repetitive constrained optimizations. The algorithm was shown to have comparable or better performance than other comparative estimators in various settings. Our aim in this paper is to further extend the algorithm to allow for variable selection in the presence of a large number of candidate markers.

Methods—We develop an alternative method to derive marker combinations to minimize the weighted sum of Ramp loss in Huang and Fong,¹ based on data from randomized trials. The new algorithm estimates treatment-selection rules by repetitively minimizing a smooth and differentiable objective function. Through the use of an L1 penalty, we expand the method to allow for feature selection and develop an algorithm based on the coordinate descent method to build the treatment-selection rule.

Results—Through extensive simulation studies, we compared performance of the proposed estimator to four existing approaches: i) a logistic regression risk modeling approach, and three other “direct optimizing” approaches including ii) the estimator in Huang and Fong,¹ iii) the weighted support vector machine; and iv) the weighted logistic regression. The proposed estimator performs comparably to that of Huang and Fong,¹ and comparably or better than other estimators. Allowing for variable selection using the proposed estimator in the presence of a large number of markers further improves treatment-selection performance. The proposed estimator is also advantageous for selecting variables relevant to treatment selection compared to L1 penalized logistic regression and weighted logistic regression. We illustrate the application of the proposed methods in host-genetics data from an HIV vaccine trial.

Conclusions—The proposed estimator is appealing considering its effectiveness and conceptual simplicity. It has significant potential to contribute to the selection and combination of biomarkers for treatment selection in clinical practice.

Keywords

biomarker; Ramp loss; total burden; treatment selection; variable selection

Introduction

In many clinical settings, intervention effects in efforts to prevent or control disease events vary across individuals. For example, breast cancer prevention research shows that the effect of estrogen therapy on postmenopausal women's risk of breast cancer varies with each individual's *FGFR2* and *MRP30* genotype.^{2,3} HIV prevention research suggests that an HIV vaccine's efficacy can be affected by various host characteristics such as the human leukocyte antigen (HLA) type and men's circumcision status.⁴ Subject-specific characteristics -- henceforth called “biomarkers” that are associated with treatment-effect heterogeneity -- can help individuals select treatment to optimize clinical outcomes.

When combining candidate markers for treatment selection, the first question to be addressed is what measure to use to quantify the treatment-selection capacity of a model. The primary goal of treatment selection is to improve prevention or control of diseases; therefore, a natural quantity to gauge the capacity of a specific rule is the rate of the targeted disease in the population as a result of treatment selection. A growing body of research in recent years has focused on developing treatment-selection rules that will minimize this quantity.^{1,5-8} Another aspect of treatment selection that has received less attention, but is as important, is the proportion of subjects in the population who will be recommended to receive the treatment and subsequent burden associated with that treatment.⁹ In practice a treatment often has disadvantages such as side effects or monetary cost. A more comprehensive measure of treatment-selection capacity should thus take into account these aspects. One example of such a measure is the “total burden” proposed by Huang and Fong.¹ This concept is defined as the sum of disease and treatment burdens in the population, where the two burdens are combined based on a pre-specified treatment-disease burden ratio following the decision theoretical framework of Vickers et al.¹⁰

In order to derive treatment-selection rules that optimize the total burden, a common strategy relies on modeling the risk of disease, conditional on the treatment and marker.^{5,11-15} Performance of rules derived in this manner relies on correct specifications of the disease risk model. An alternative strategy that is more robust to model misspecification is to minimize an unbiased estimate of total burden within a pre-specified class of treatment-selection rules.^{1,6-8,16,17} This can be framed as an equivalent problem of minimizing a weighted sum of 0-1 loss function, which is computationally challenging to solve due to the non-smoothness and non-convexity of the 0-1 loss. Researchers have therefore developed methods that approximate the 0-1 loss with various surrogate loss functions, e.g. the convex hinge loss function underlying the support vector machine method⁸ and the binomial deviance loss function underlying the logistic regression method.¹⁷ Huang and Fong¹

proposed a penalized estimator that minimizes a weighted sum of the Ramp loss function, which approximates the 0-1 loss more closely compared to the hinge loss, and showed that using a more precise approximation can often improve performance. They developed a “difference of convex functions algorithm” (DCA), which consists of repetitively solving a constrained convex optimization problem.¹

In this paper, we develop an alternative DCA algorithm to minimize the penalized weighted sum of Ramp loss adopted by Huang and Fong.¹ In particular, in each step of the DCA, we approximate the objective function to be minimized by a smooth convex function. The conceptual simplicity of minimizing a smooth convex objective function is appealing; it involves standard statistical elements familiar to general biostatistics researchers. Most importantly, we extend the algorithm to accommodate the need of feature selection when combining multiple markers through the use of an L1 penalty. We compare finite-sample performance of this method with other existing methods through simulation studies and in a real example from an HIV vaccine trial.

Methods

We consider data from a two-arm randomized trial. Let Y indicate an unfavorable binary disease outcome taking values 0 and 1, where 0 stands for non-diseased and 1 stands for diseased. Let T indicate the treatment assignment a subject received in the trial, with values 0 and 1 indicating untreated and treated, respectively. Let X indicate a set of biomarkers with dimension p . Our objective is to derive a marker-based binary treatment-selection rule $A(X)$ using the trial data for making treatment recommendations in the future such that subjects with biomarker value $X=x$ will be recommended to receive the treatment if $A(x)=1$ and to not receive the treatment if $A(x)=0$.

Let $Y(1)$ and $Y(0)$ indicate the potential outcomes of Y if a subject were to receive or not receive the treatment. If every subject in the population follows the treatment recommendation given by $A(X)$, the average population disease rate will equal $E_{A(X)}(Y) = E[Y(0) \times I\{A(X)=0\}] + E[Y(1) \times I\{A(X)=1\}]$.^{5,6} Moreover, following the decision-theoretic framework,¹⁰ suppose the burden of treating a subject due to side effects or monetary cost is δ times the burden per disease event, where δ is a pre-specified burden ratio quantifying how significant the treatment burden is relative to the burden of having the targeted disease. Then the sum of disease and treatment burdens, i.e. the total burden measured on the unit of burden per disease event, equals $E_{A(X)}(Y) + P\{A(X)=1\} \times \delta$. An example of the choice of δ was provided by Vickers et al.¹⁰ where δ was assigned a value of 5% through a patient survey for treating breast cancer with adjuvant chemotherapy. This corresponds to assuming the harm of death due to breast cancer is 20 times the harm of having adjuvant chemotherapy.

With data from a randomized trial, total burden can be represented as $E(Y|T=0) - E[\{ (X) - \delta \} \times A(X)]$ (see Supplementary Appendix A), where $(X) = P(Y=1 | X, T=0) - P(Y=1 | X, T=1)$ is the difference in disease risk conditional on the marker between untreated and treated. This indicates that the optimal rule that minimizes the total burden is to treat (i.e. $A(X)=1$) if $(X) > \delta$ and to not treat (i.e. $A(X)=0$) otherwise. Therefore, regression-based methods can be

used to model the disease risk and estimate the risk difference $E\{f(X) - \delta\}$ to derive the rule. This method is efficient when the risk model is correctly specified but could otherwise have suboptimal performance.¹⁸ An alternative strategy, more robust to model misspecification, is to directly minimize an unbiased estimate of the criterion function among a pre-specified class of rules.^{1,6-8,16,17} Consider $A(X) = I\{f(X) > 0\}$ for $f(X)$ a smooth function of X . Minimizing the total burden $E(Y|T=0) - E\{(X) - \delta\} \times I\{f(X) > 0\}$ is equivalent to minimizing $E\{(X) - \delta\} \times I\{f(X) > 0\}$. A consistent estimate of the latter quantity can be written as a weighted sum of 0-1 loss:

$$\sum W_i \times I\{f(X) \leq 0\}, \quad (1)$$

where W_i is a subject-specific weight that is a consistent estimate of $(X) - \delta$ [1]. Therefore, the estimation of $f(X)$ to minimize the total burden can be carried out as a problem of minimizing (1).

Two principal choices of the weights W are the outcome-weights and the double-robustness weights.^{1,6-8} First, in a randomized trial, total burden can be written as $E[Y \times I\{f(X) > 0\} | T=0] + E[Y \times I\{f(X) > 0\} | T=1] + \delta \times E\{I\{f(X) > 0\}\} = E[Y \times T \times I\{f(X) > 0\}] / P(T=1) + E[Y \times (1-T) \times I\{f(X) > 0\}] / P(T=0) + \delta \times P\{f(X) > 0\}$. The empirical estimate of this quantity equals

$$\hat{\theta}_1 = \sum_{i=1}^n Y_i \times T_i \times I\{f(X_i) > 0\} / n_1 + Y_i \times (1 - T_i) \times I\{f(X_i) \leq 0\} / n_0 + \delta \times I\{f(X_i) > 0\} / n,$$

where n_0, n_1 are numbers of untreated and treated, respectively, and $n = n_0 + n_1$. Minimizing $\hat{\theta}_1$ is equivalent to minimizing (1) with subject-specific weights

$W_i = - \left\{ \frac{Y_i \times T_i}{n_1/n} - \frac{Y_i \times (1 - T_i)}{n_0/n} + \delta \right\}$ (see Supplementary Appendix B). Second, a doubly-robust estimator $\hat{\theta}_2$ of total burden can be constructed, which augments the empirical estimator $\hat{\theta}_1$ with a term involving the risk of Y given X and T :

$$\hat{\theta}_2 = \hat{\theta}_1 - \sum_{i=1}^n \frac{(1 - T_i) \times I\{f(X_i) \leq 0\} + T_i \times I\{f(X_i) > 0\} - P[I\{f(X_i) > 0\} = T | X_i]}{P[I\{f(X_i) > 0\} = T | X_i]} \times \hat{m}(X_i; f),$$

where $P[I\{f(X) > 0\} = T | X] = \pi \times I\{f(X) > 0\} + (1 - \pi) \times I\{f(X) \leq 0\}$ for $\pi = P(T=1)$, $\hat{m}(X_i; f) = P(\hat{Y} = 1 | X_i, T = 0) \times I\{f(X_i) \leq 0\} + P(\hat{Y} = 1 | X_i, T = 1) \times I\{f(X_i) > 0\}$, with $P(\hat{Y} = 1 | X, T)$ estimated based on a working model for disease risk.^{6,19} This estimator has the double-robustness property in that it is consistent if either $P(T=1|X)$ or the risk model is correctly specified. In a randomized trial, since $P(T=1|X)$ is known, consistency of the estimator is always achievable; using the augmented term has the added advantage of increasing asymptotic efficiency if the risk model is correctly specified.⁶ Minimizing the doubly-robust estimator is equivalent to minimizing (1) with subject-specific weights

$$W_i = - \left[\frac{Y_i T_i}{n_1} - \frac{Y_i (1 - T_i)}{n_0} + \frac{\pi - T_i}{1 - \pi} \times \frac{\hat{P}(Y=1|X_i, T=0) + \hat{P}(Y=1|X_i, T=1)}{n} + \frac{\delta}{n} \right]$$

(see Supplementary Appendix C).¹

Minimizing the weighted sum of 0-1 loss (1) is computationally challenging due to the non-smoothness and non-convexity of 0-1 loss. Zhao et al.⁸ proposed using the support vector machine method for minimization, which essentially approximates the 0-1 loss with a convex hinge loss function (Figure 1a). Huang and Fong¹ adopted a Ramp loss function $h(x)$, which approximates the 0-1 loss function more closely (Figure 1a), where $h(x)=1$ if $x \leq -1/2$, $h(x)=1/2-x$ if $-1/2 < x < 1/2$, and $h(x)=0$ if $x \geq 1/2$. In this paper, we introduce a penalized estimator for deriving a linear marker combination $f(X)=b+X^T\beta$ based on the Ramp loss approximation.

The weighted smooth ramp estimator

We propose to solve for b and β by minimizing a weighted sum of Ramp loss with an L2 penalty: $\sum |W_i| \times h[\text{sgn}(W_i) \times (b + X_i^T \beta)] / n + 0.5 \lambda_n \|\beta\|^2$, where λ_n is a tuning parameter, and $\text{sgn}(w)$ indicates the sign of w . The Ramp loss function can be written as the difference of two convex functions (Figure 1b), $h(x) = h_1(x) - h_2(x)$ with $h_1(x) = (1/2 - x)_+$ and $h_2(x) = (-1/2 - x)_+$. Therefore, the “difference of convex functions algorithm” (DCA) can be used to estimate the marker combination.

Let $\eta_i = \text{sgn}(w_i) \times f(x_i) = \text{sgn}(w_i) \times (b + x_i^T \beta)$, the minimization is carried out in three steps:

1. Start with an initial estimate for η and call it η^0 .

2. Solve $\hat{\beta} = \arg\min_{\beta} \frac{1}{n} \sum_{i=1}^n |w_i| \{h_1(\eta_i) - \hat{h}_2(\eta_i, \eta_i^0)\} + \frac{1}{2} \lambda_n \|\beta\|_2^2$, 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 $u = -1/2$.

We propose to solve this convex optimization problem by approximating the hinge

function $h_1(u)$ with a smooth function, $S(u) = 0.5 - \varepsilon \log \frac{1 + \exp(0.5/\varepsilon)}{1 + \exp\{(0.5 - u)/\varepsilon\}}$, as shown in Figure 1 b. This approximation has been previously used in a different problem setting for maximizing the partial area under the Receiver Operating Characteristic curve.²⁰ For $\varepsilon \rightarrow 0$, the difference between $S(u)$ and $h_1(u)$ converges to zero. Replacing $h_1(u)$ with $S(u)$ for a small value of ε , the objective function minimized in Step 2 then becomes a smooth convex function whose first and second derivatives exist everywhere. Standard statistical algorithms for minimizing convex and differentiable objective functions can be used to solve for β and b . In our simulation and data example, we use a quasi-newton method, the BFGS, implemented in R optim function²¹ for minimization.

3. Compute η^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.

Note that a similar three-step approach was used earlier in Huang and Fong,¹ where they parameterized Step 2 into a constrained optimization problem and solved its dual problem via quadratic programming methods. Here we adopt the smooth convex approximation to the objective function such that standard optimization algorithms accessible to general biostatistical researchers can be used with little sacrifice of performance.

Feature selection

We further extend the penalized smoothed Ramp estimator to allow for feature selection when combining markers. In practice, when a large number of candidate biomarkers are available to combine, feature selection is desirable for several reasons: (i) biomarkers can be costly to measure. It is cost-effective to generate a model based on a parsimonious set of markers that are most effective for treatment-selection purposes; (ii) it is beneficial to select a small set of useful biomarkers to interpret and understand underlying biological mechanisms; and (iii) estimated models based on many covariates can be unstable in finite samples with suboptimal performance ---feature selection helps alleviate this issue.

To incorporate feature selection, we propose to minimize the weighted sum of Ramp loss function with an L1 penalty. That is, we solve for

$\sum |W_i| h\{\text{sgn}(W_i) \times (b + X_i^T \beta)\} / n + 0.5 \lambda_n \sum_j |\beta_j|$. The minimization can again be carried out in three steps as above, and Steps 1 and 3 are the same. In Step 2, the objective becomes

estimating $\hat{\beta} = \arg\min_{\beta} \frac{1}{n} \sum_{i=1}^n |w_i| \{h_1(\eta_i) - \hat{h}_2(\eta_i, \eta_i^0)\} + \frac{1}{2} \lambda_n \sum_j |\beta_j|$. Replacing $h_1(u)$ with $S(u)$, we solve for β using standard algorithms for minimizing L1 penalized smooth and differentiable function such as the coordinate descent algorithm.²² We update parameters individually. In each step, we solve for a particular parameter b or $\beta_j, j \in \{1, \dots, p\}$ to

minimize $\frac{1}{n} \sum_{i=1}^n |w_i| \{S(X_i^T \beta) - \hat{h}_2(X_i^T \beta, \eta_i^0)\} + \frac{1}{2} \lambda_n \sum_j |\beta_j| \propto \psi(\beta_j) + \frac{1}{2} \lambda_n |\beta_j|$, while other

parameters are held constant, where $\psi(\beta_j) \equiv \frac{1}{n} \sum_{i=1}^n |w_i| \{S(X_i^T \beta) - \hat{h}_2(X_i^T \beta, \eta_i^0)\}$. For

example, to estimate β_j we compute $\tilde{\beta}_j^a = \arg\min_{\beta_j} \psi(\beta_j) + \frac{1}{2} \lambda_n \beta_j$ and

$\tilde{\beta}_j^b = \arg\min_{\beta_j} \psi(\beta_j) - \frac{1}{2} \lambda_n \beta_j$, then compare $\psi(\tilde{\beta}_j^a) + \frac{1}{2} \lambda_n \tilde{\beta}_j^a$, $\psi(\tilde{\beta}_j^b) + \frac{1}{2} \lambda_n \tilde{\beta}_j^b$ and $\Psi(0)$. If

$\psi(\tilde{\beta}_j^a) + \frac{1}{2} \lambda_n \tilde{\beta}_j^a$ is the smallest among the three and $\tilde{\beta}_j^a > 0$, then we update β_j with $\tilde{\beta}_j^a$; If

$\psi(\tilde{\beta}_j^b) - \frac{1}{2} \lambda_n \tilde{\beta}_j^b$ is the smallest among the three and $\tilde{\beta}_j^b < 0$, then we update β_j with $\tilde{\beta}_j^b$; otherwise β_j is updated to zero. R code for implementing the algorithm can be found in <http://research.fhcrc.org/huang/en.html>.

Simulation studies

We conduct simulation studies to investigate performance of the proposed methods. Without loss of generality, here we focus on settings where treatment cost is unremarkable compared to disease cost such that cost ratio δ can be set to zero. The objective of treatment selection is thus to minimize the average population disease rate, which equals total burden.

We consider data from 300 subjects, 1:1 randomized to untreated and treated groups. Overall 50 biomarkers $X_1, \dots, X_{50}$ are measured, among which X_1 and X_2 contribute to the disease risk model and the remainder are noisy variables. We assess the proposed weighted smooth RAMP estimator (W-SRAMP) for deriving linear marker combinations $f(X) = b + X^T \beta$ for treatment selection. We explore ε from 0.001 to 0.5 and replace $\log \{1 + \exp((0.5-u)/\varepsilon)\}$ in $S(u)$ with $(0.5-u)/\varepsilon$ when $(0.5-u)/\varepsilon \geq 700$ to guard against machine overflow.

We also study four comparative estimators: i) the estimator based on a linear logistic regression model (LOGISTIC) that includes main effect for treatment and each individual marker, and the interaction between treatment and individual marker; and ii) the weighted RAMP estimator (W-RAMP) proposed in Huang and Fong¹ for minimizing the penalized weighted sum of Ramp loss function. Furthermore, since

$$\begin{aligned} \arg\min_f \sum_{i=1}^n W_i I\{f(X_i) \leq 0\} &= \arg\min \sum_{i=1}^n |W_i| I(W_i > 0) I\{f(X_i) \leq 0\} - |W_i| I(W_i \leq 0) I\{f(X_i) \leq 0\} \\ &= \arg\min \sum_{i=1}^n |W_i| I(W_i > 0) I\{f(X_i) \leq 0\} + |W_i| I(W_i \leq 0) I\{f(X_i) > 0\} \\ &= \arg\min \sum_{i=1}^n |W_i| I\{\text{sgn}(W_i) \times f(X_i) \leq 0\} = \arg\min \sum_{i=1}^n |W_i| I[\text{sgn}(W_i) \neq \text{sgn}\{f(X_i)\}], \end{aligned}$$

minimizing the weighted sum of 0-1 loss (1) can be translated into a problem of minimizing weighted classification error,^{1,6-8,17} where the goal is to derive a rule $f(X_i)$ to predict the binary outcome characterized by $\text{sgn}(W_i)$, with subject-specific weights $|W_i|$. The third and fourth comparative estimators we investigate are: iii) the weighted support vector machine estimator (W-SVM);^{1,8} and iv) the weighted logistic regression estimator (W-LOGIS),¹⁷ both regressing $\text{sgn}(W_i)$ on X_i with weights $|W_i|$. The five-fold cross-validation procedure is used to select the tuning parameter λ_n in W-SRAMP, W-RAMP, and W-SVM within a range of 0 to 5 after $\sum_{i=1}^n |W_i|/n$ is scaled to be 1.

Moreover, we compare the following estimators that incorporate feature selection into the derivation of marker combinations: i) the linear logistic regression risk model described above + L1-norm penalty (LOGISTIC-L1); ii) the weighted logistic regression estimator + L1 norm penalty (W-LOGIS-L1); and iii) the proposed estimator minimizing weighted smooth RAMP loss + L1-norm penalty (W-SRAMP-L1). Five-fold cross-validation is used for selecting the tuning parameter λ_n in W-SRAMP-L1 within a range of 0 to 0.5 after

$\sum_{i=1}^n |W_i|/n$ is scaled to be 1. LOGISTIC-L1 and W-LOGIS-L1 are implemented with the `cv.glmnet` function in the R package `glmnet` under the default parameter setting;²³ after variable selection, the set of variables with nonzero coefficients are refit to the training data to generate the marker model in order to alleviate the bias in estimating model coefficients as suggested by Meinshausen.²⁴

All estimators are evaluated in the following five scenarios with marker data $X_1, \dots, X_{50}$ independent of each other. Individual markers follow a standard normal distribution in all scenarios except for scenario 5.

In scenario 1, data are simulated from a linear logistic model with linear marker main effect and interaction with treatment: $\text{logit } P(Y=1|X, T) = -1.5 + 0.5X_1 - 0.5X_2 + T \times (-1 - 1.5X_1 - 1.5X_2)$. Thus the working model assumed in LOGISTIC is correctly specified (which is not true in other scenarios), and the optimal treatment-selection rule is based on linear marker combinations.

In scenario 2, data are simulated from a logistic regression model

$\text{logit } P(Y=1|X, T) = -1.5 + 0.5X_1 - 0.5X_2 - 0.5X_1^2 + 0.5X_2^2 + T \times (-0.35 - 1.5X_1 + 1.5X_2)$. The markers' main effects involve quadratic terms, while their interactions with treatment involve only linear terms such that the optimal treatment-selection rule is again based on linear marker combinations.

For scenarios 3-5, we investigate performance of various estimators of best linear marker combinations when the overall optimal marker combination is not contained in the linear class.

In Scenario 3, data are simulated from a logistic regression model

$\text{logit } P(Y=1|X, T) = -1.4 + 0.5X_1 - 0.5X_2 + T \times (-1 - 1.5X_1 - 1.5X_2 + 0.5X_1^2 - 0.5X_2^2)$. The markers' main effects are linear, but their interactions with treatment involve quadratic terms.

In Scenario 4, data are simulated from a logistic regression model

$\text{logit } P(Y=1|X, T) = -1.5 + 0.5X_1 - 0.5X_2 - 0.5X_1^2 + 0.5X_2^2 + T \times [2.5 - 4 \times I(X_1 > 0) - 4 \times I(X_2 > 0) - 2 \times I(X_1 > 0) \times I(X_2 > 0)]$. Both main effect and interaction terms are nonlinear.

In Scenario 5, data are generated from a mixture model. With 90% probability, $X_1, \dots, X_{50}$ are independent standard normal random variables and the risk of Y given X and T follows

$\text{logit } P(Y=1|X, T) = -2.8 + 0.2X_1 - 0.2X_2 + 0.5X_1^2 + 0.5X_2^2 + T \times (-1 - X_1 - X_2 + 0.5X_1^2 + 0.5X_2^2)$; with 10% probability, $X_1, \dots, X_{50}$ are independent normal random variables with mean 2 for X_1 and 0 for others, and variance 1.5 for X_1 and X_2 and 1 for others, the risk of Y given X and T follows

$\text{logit } P(Y=1|X, T) = -2.8 + 0.2X_1 - 0.2X_2 + 0.5X_1^2 + 0.5X_2^2 - 0.5X_1X_2 + T \times (-3 - 0.5X_1^2 - 0.5X_2^2)$.

In all of these scenarios, disease prevalence is higher among untreated than treated, so the optimal strategy absent the marker information is to treat everyone in the population. A summary of the five scenarios is presented in Table 1.

For each scenario, 500 Monte-Carlo simulations are performed. We estimate the treatment-selection rule using a simulated training data of 300 individuals, based on tuning parameters selected specifically for this training data using procedures described above. The estimated rules are then applied to an independently simulated test data of 10,000 individuals, where $\hat{\theta}$,

the disease rate as a result of treatment selection, is estimated. The average of $\hat{\theta}$ over 500 simulations is calculated.

Table 2 shows the performance of various estimators of marker combinations using the complete set of 50 markers without variable selection, based on either the outcome-weights or the double-robustness weights. For the latter weights, the augmented term is derived based on the linear logistic model used in LOGISTIC. For ease of presentation, we present the benefit of treatment selection, which is defined as the reduction in the disease rate as a result of marker-based treatment selection related to the default strategy of treating all. Overall, W-RAMP and W-SRAMP perform best in all settings, followed by W-SVM and then LOGISTIC and W-LOGIS estimators. By approximating $h_1(u)$, W-SRAMP can have a small decrease in performance compared to W-RAMP in some settings, but still significantly outperforms LOGISTIC and W-LOGIS. Using double-robustness weights in general improves performance of W-SVM, and sometimes of W-LOGIS as well, but has an unremarkable impact on performance of W-SRAMP or W-RAMP. Results presented for W-RAMP are based on $\varepsilon=0.5$; with a large number of noisy markers to combine, using smaller ε can lead to slight reduction of performance, but the change is minimal relative to the comparison with other estimators.

Table 3 presents performance of estimators with feature selection, i.e., W-SRAMP-L1, LOGISTIC-L1, W-LOGIS-L1, using outcome weights or double-robustness weights. The augmented term for the latter weights is derived using the model from LOGISTIC-L1. Performances of marker combinations derived after variable selection have apparent improvement over marker combinations derived without variable selection as presented in Table 2. W-SRAMP-L1 performs best among the three in most scenarios except in scenario 1, where linear logistic model is true and LOGISTIC-L1 performs better than W-SRAMP-L1 with outcome-weights. Double-robustness weights in general improve performance of estimators over outcome-weights. Results of W-SRAMP-L1 are fairly insensitive to choice of ε .

Table 4 shows sensitivity and specificity of LOGISTIC-L1, W-LOGIS-L1 and W-SRAMP-L1, where sensitivity is defined as the average probability that a useful marker (X_1, X_2) is selected, and specificity is defined as the average probability that a noisy marker ($X_3, \dots, X_{50}$) is not selected. For all five scenarios, the three estimators have very good specificity (>95%). W-SRAMP-L1 has appreciably improved sensitivity compared to W-LOGIS-L1 or LOGISTIC-L1. Double-robustness weights in general lead to estimators with better sensitivity compared to outcome-weights.

Data example

We illustrate the proposed methods for selecting and combining markers with an example from the RV144 Thailand HIV vaccine trial, the first HIV vaccine trial that demonstrates a positive vaccine efficacy in preventing HIV infection, which included 16,402 participants aged 18-30 who were 1:1 randomized into a vaccine and a placebo arm.²⁵ A followup RV144 host-genetic study was conducted to investigate the effect of genotypes of Fc receptor gene on vaccine efficacy, where 190 single nucleotide polymorphisms (SNPs)

covering five Fc- γ receptor genes and 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), 28 tags SNPs passing quality control based on Hardy-Weinberg equilibrium were selected by Li et al.²⁶ to study association with vaccine efficacy, each categorized into a binary variable.

Here we explore combinations of the 28 SNPs, together with an individual's age, gender, and baseline behavioral risk to recommend the use of vaccine. As a prevention measure, vaccines typically have low cost related to treatment, thus we set $\delta=0$ and compare all estimators of linear marker combinations as presented in simulation studies with respect to disease rates resulted from treatment selection.

A random cross-validation procedure is performed for the computation of disease rate resulting from treatment selection to correct for over-fitting 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 the disease rate. The procedure is repeated 100 times with the average disease rate computed.

Table 5 shows estimated performance of different estimators. The strategy of treating none and treating all lead to an estimated HIV infection rate of 8.54 and 6.44 per 1000 persons respectively, consistent with the positive vaccine efficacy we observed in the RV144 trial. Among methods without feature selection, estimators that minimize the criterion functions directly (i.e. W-LOGIT, W-SVM, W-RAMP, and W-SRAMP) perform better compared to the risk-model based estimator (LOGISTIC). However, none of the methods reduce population disease rate compared to treating all. Feature selection in general leads to better rules. The rule that performs best is W-SRAMP-L1 with double-robustness weights, which recommends 96.5% subjects to receive vaccination and leads to a reduced HIV infection rate of 6.34 per 1000 persons compared to 6.44 per 1000 persons for the treating-all strategy.

Summary and conclusion

In this paper, we consider utilizing biomarkers to help individuals make treatment decisions. When the goal is to develop treatment-selection rules to minimize population burdens due to disease and treatment, the strategy to directly minimize the targeted criterion function within a pre-specified class of rules, namely the “policy search methods”,²⁷ requires fewer model assumptions and thus is more robust compared to the standard strategy of modeling the entire disease risk.

The criterion function optimized in policy search methods involves a non-smooth 0-1 loss, which requires replacement by a surrogate loss function for computational efficiency. In this paper, we propose an algorithm to derive linear marker combination for treatment selection that minimizes the weighted Ramp loss function proposed in Huang and Fong.¹ In each

iteration, the proposed algorithm requires minimizing a smooth and differentiable objective function, which presents more ‘standard’ questions that are accessible to general biostatistics fields relative to the algorithm proposed earlier by Huang and Fong¹ based on constrained optimization. The idea of smooth approximation is applicable to other problem settings where there is a need to optimize a non-smooth function and it is relatively easy to implement. Most importantly, we extend the algorithm to incorporate feature selection when a large number of candidate markers are present. Through comparison with alternative estimators based on a logistic regression risk model or a weighted logistic regression policy search method, the proposed estimator achieves comparable or better performance both in minimizing the total disease and treatment burden and in selecting relevant biomarkers. Therefore, the method has a strong likelihood of complementing existing methods and contributing to policy search in clinical practice.

Supplementary Appendix

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

The author thanks Drs. Dan Geraghty, Sue Li and Peter Gilbert from the Fred Hutchinson Cancer Research Center for providing the host-genetics data, and thanks Dr. Youyi Fong for helpful discussions. The work was supported by NIH grant R01 GM106177-01.

References

1. Huang Y, Fong Y. Identifying optimal biomarker combinations for treatment selection via a robust kernel method. *Biometrics*. 2014; 70:891–901. [PubMed: 25124089]
2. Prentice RL, Huang Y, Hinds DA, et al. Variation in the FGFR2 gene and the effects of postmenopausal hormone therapy on invasive breast cancer. *Cancer Epidemiol Biomarkers Prev*. 2009; 18:3079–3085. [PubMed: 19861516]
3. Huang Y, Ballinger DG, Dai JY, et al. Genetic variants in the MRPS30 region and postmenopausal breast cancer risk. *Genome Med*. 2011; 3:42. [PubMed: 21702935]
4. Buchbinder SP, Mehrotra DV, Duerr A, et al. Efficacy assessment of a cell-mediated immunity HIV-1 vaccine (the Step Study): a double-blind, randomised, placebo-controlled, test-of-concept trial. *Lancet*. 2008; 372:1881–1893. [PubMed: 19012954]
5. Song X, Pepe MS. Evaluating markers for selecting a patient's treatment. *Biometrics*. 2004; 60:874–883. [PubMed: 15606407]
6. Zhang B, Tsiatis AA, Laber E, et al. A robust method for estimating optimal treatment regimes. *Biometrics*. 2012; 68:1010–1018. [PubMed: 22550953]
7. Zhang B, Tsiatis AA, Davidian M, et al. Estimating optimal treatment regimes from a classification perspective. *Stat*. 2012; 1:103–114. [PubMed: 23645940]
8. Zhao Y, Zeng D, Rush AJ, et al. Estimating individualized treatment rules using outcome weighted learning. *J Am Stat Assoc*. 2012; 107:1106–1118. [PubMed: 23630406]
9. Janes H, Pepe MS, Bossuyt PM, et al. Measuring the performance of markers for guiding treatment decisions. *Annals Intern Med*. 2011; 154:253–259.
10. Vickers A, Kattan M, Sargent D. Method for evaluating prediction models that apply the results of randomized trials to individual patients. *Trials*. 2007; 8:14. [PubMed: 17550609]
11. Janes H, Pepe M, Huang Y. A framework for evaluating markers used to select patient treatment. *Med Decis Making*. 2013; 34:159–167. [PubMed: 23811760]
12. Foster JC, Taylor JMG, Ruberg SJ. Subgroup identification from randomized clinical trial data. *Stat Med*. 2011; 30:2867–2880. [PubMed: 21815180]

13. Lu W, Zhang HH, Zeng D. Variable selection for optimal treatment decision. *Stat Methods Med Res.* 2013; 22:493–504. [PubMed: 22116341]
14. McKeague IW, Qian M. Estimation of treatment policies based on functional predictors. *Stat Sin.* 2014; 24:1461–1485. [PubMed: 25165416]
15. Kang C, Janes H, Huang Y. Combining biomarkers to optimize patient treatment recommendations. *Biometrics.* May 30.2014 Epub ahead of print. 10.1111/biom.12191
16. Orellana L, Rotnitzky A, Robins JM. Dynamic regime marginal structural mean models for estimation of optimal dynamic treatment regimes, part i: main content. *Int J Biostat.* 2010; 6 Article 8.
17. Huang, Y.; Laber, EB.; Janes, H. Characterizing expected benefits of biomarkers in treatment selection. Univ of Washington Working Paper. 2014. <http://biostats.bepress.com/uwbiostat/paper396/>
18. Qian M, Murphy SA. Performance guarantees for individualized treatment rules. *Ann Stat.* 2011; 39:1180–1210. [PubMed: 21666835]
19. Robins JM, Rotnitzky A, Zhao LP. Estimation of regression coefficients when some regressors are not always observed. *J Am Stat Assoc.* 1994; 89:846–866.
20. Wang Z, Chang YCI. Marker selection via maximizing the partial area under the ROC curve of linear risk scores. *Biostatistics.* 2011; 12:369–385. [PubMed: 20729218]
21. Nash, JC. Compact numerical methods for computers. Linear algebra and function minimisation. 2nd. Bristol and New York: Adam Hilger; 1990.
22. Friedman J, Hastie T, Höfling H, et al. Pathwise coordinate optimization. *Ann Appl Stat.* 2007; 1:302–332.
23. Friedman J, Hastie T, Tibshirani R. Regularization paths for generalized linear models via coordinate descent. *J Stat Softw.* 2010; 33:1–22. [PubMed: 20808728]
24. Meinshausen N. Relaxed lasso. *Comput Stat Data Anal.* 2007; 52:374–393.
25. Rerks-Ngarm S, Pitisuttithum P, Nitayaphan S, et al. Vaccination with ALVAC and AIDSVAX to prevent HIV-1 infection in Thailand. *N Engl J Med.* 2009; 361:2209–2220. [PubMed: 19843557]
26. Li SS, Gilbert PB, Tomaras GD, et al. FCGR2C polymorphisms associate with HIV-1 vaccine protection in RV144 trial. *J Clin Invest.* 2014; 124:3879–3890. [PubMed: 25105367]
27. Laber EB, Tsiatis AA, Davidian M, et al. Discussion of “combining biomarkers to optimize patient treatment recommendation”. *Biometrics.* Jun 2.2014 Epub ahead of print. 10.1111/biom.12187

Figure 1.

(a) Different approximations of the 0-1 loss function. (b) Representing the ramp loss $h(u)$ as the difference of two convex functions $h_1(u)$ and $h_2(u)$, and the approximation of $h_1(u)$ with $S_\epsilon(u)$.

Table 1

Summary of simulation settings.

Setting	Disease risk model: $\text{logit } P(Y=1 X,T)$	Distribution of $X_1, \dots, X_{60}$	$P(Y=1 T=0)$	$P(Y=1 T=1)$
1	$-1.5+0.5X_1-0.5X_2+T \times (-1-1.5X_1-1.5X_2)$	$N(0,1)$	0.204	0.189
2	$-1.5+0.5X_1-0.5X_2-0.5X_1^2+0.5X_2^2+T \times (-0.35-1.5X_1+1.5X_2)$	$N(0,1)$	0.225	0.208
3	$-1.4+0.5X_1-0.5X_2+T \times (-1-1.5X_1-1.5X_2+0.5X_1^2-0.5X_2^2)$	$N(0,1)$	0.219	0.200
4	$-1.5+0.5X_1-0.5X_2-0.5X_1^2+0.5X_2^2+T \times (2.5-4 \times I(X_1>0)-4 \times I(X_2>0)+2 \times I(X_1>0) \times I(X_2>0))$	$N(0,1)$	0.225	0.214
5	With $p=0.9$ $-2.8+0.2X_1-0.2X_2+0.5X_1^2+0.5X_2^2+T \times (-1-X_1-X_2+0.5X_1^2+0.5X_2^2)$	$N(0,1)$		
	With $p=0.1$ $-2.8+0.2X_1-0.2X_2+0.5X_1^2+0.5X_2^2-0.5X_1X_2+T \times (-3-0.5X_1^2-0.5X_2^2)$	$X_1 \sim N(2,1.5), X_2 \sim N(0,1.5)$, others $N(0,1)$	0.213	0.208

Table 2
Reduction in disease rate using marker-based rule compared to the strategy of treating all.

Scenario	Weights	W-SRAMP [#]	Logistic	W-Logit	W-RAMP	W-SVM
1	OC ^a	2.92 (0.06)	2.65 (0.07) [*]	2.56 (0.05) [*]	3.05 (0.05) [*]	2.58 (0.06) [*]
	DR ^b	3.03 (0.05)	2.65 (0.07) [*]	1.91 (0.07) [*]	3.02 (0.05)	2.87 (0.05) [*]
2	OC	4.23 (0.07)	3.50 (0.07) [*]	3.04 (0.06) [*]	4.23 (0.08)	3.90 (0.08) [*]
	DR	4.19 (0.06)	3.50 (0.07) [*]	3.58 (0.07) [*]	4.22 (0.07)	4.17 (0.06)
3	OC	2.96 (0.06)	2.59 (0.06) [*]	1.93 (0.06) [*]	3.06 (0.07) [*]	2.48 (0.07) [*]
	DR	3.02 (0.06)	2.59 (0.06) [*]	1.89 (0.06) [*]	3.11 (0.06) [*]	2.85 (0.05) [*]
4	OC	2.22 (0.05)	2.25 (0.06)	2.25 (0.05)	2.34 (0.05) [*]	1.89 (0.06) [*]
	DR	2.32 (0.05)	2.25 (0.06)	1.23 (0.05) [*]	2.40 (0.05) [*]	2.08 (0.04) [*]
5	OC	3.33 (0.06)	2.60 (0.05) [*]	1.62 (0.04) [*]	3.30 (0.07)	2.97 (0.07) [*]
	DR	3.28 (0.06)	2.60 (0.05) [*]	2.64 (0.05) [*]	3.29 (0.06)	3.36 (0.06)

OC^a: outcome-weight

DR^b: double-robustness weight

W-SRAMP[#]: with $\epsilon=0.5$

^{*} indicates p-value for comparing the estimator with W-SRAMP based on two-sided paired-t test is less than 0.05

Table 3

Reduction in disease rate using marker-based rule compared to the strategy of treating all.

Scenario	Weights	W-SRAMP-L1 [#]	Logistic-L1	W-Logit-L1
1	OC ^a	4.51 (0.11)	4.99 (0.12) [*]	3.97 (0.14) [*]
	DR ^b	5.69 (0.09)	4.99 (0.12) [*]	5.38 (0.12) [*]
2	OC	6.61 (0.12)	3.36 (0.19) [*]	5.72 (0.16) [*]
	DR	7.55 (0.09)	3.36 (0.19) [*]	7.51 (0.11)
3	OC	5.29 (0.12)	4.60 (0.17) [*]	3.91 (0.16) [*]
	DR	6.00 (0.09)	4.60 (0.17) [*]	5.63 (0.12) [*]
4	OC	4.98 (0.14)	3.25 (0.13) [*]	3.90 (0.16) [*]
	DR	5.82 (0.12)	3.25 (0.13) [*]	4.16 (0.15) [*]
5	OC	4.06 (0.11)	2.23 (0.14) [*]	2.21 (0.12) [*]
	DR	4.71 (0.09)	2.23 (0.14) [*]	4.60 (0.12)

OC^a: outcome-weight

DR^b: double-robustness weight

W-SRAMP-L1[#]: with $\varepsilon=0.1$

^{*} indicates p-value for comparing the estimator with W-SRAMP-L1 based on two-sided paired-t test is less than 0.05.

Average sensitivity and specificity of various methods for selecting biomarkers relevant to treatment selection.

Table 4

Method	Weights	Scenario1		Scenario2		Scenario3		Scenario4		Scenario5	
		SEN	SPE	SEN	SPE	SEN	SPE	SEN	SPE	SEN	SPE
Logistic-L1		60.5	99.5	27.7	99.7	51.7	99.7	47.9	99.8	30.6	99.6
W-Logit-L1	OC ^a	57.3	97.5	64.4	98.0	40.9	98.3	44.3	97.7	28.6	98.2
W-Logit-L1	DR ^b	71.4	97.7	73.1	98.3	50.0	98.3	52.4	97.8	48.3	98.4
W-SRAMP-L1 [#]	OC	63.5	96.3	68.0	97.2	53.9	96.5	57.5	96.2	56.5	95.2
W-SRAMP-L1 [#]	DR	81.4	96.8	80.6	97.6	60.2	97.0	66.4	96.8	61.1	96.0

OC^a: outcome-weight

DR^b: double-robustness weight

W-SRAMP-L1[#]: with $\varepsilon=0.1$

Table 5

Mean cross-validated estimates of disease prevalence $\times 1000$ as the result of different treatment strategy in the RV144 host-genetics study.

Strategy	Weights	
	OC ^a	DR ^b
Treating All	6.44	
Treating None	8.54	
Logistic	8.37	
W-Logit	6.81	6.52
W-RAMP	6.79	6.73
W-SVM	6.89	6.59
W-SRAMP [#]	6.77	6.53
Logistic-L1	6.44	
W-Logit-L1	6.45	6.37
W-SRAMP-L1 [#]	6.73	6.34

OC^a: outcome-weight

DR^b: double-robustness weight

W-SRAMP[#] and W-SRAMP-L1[#]: with $\varepsilon=0.1$.