Published in final edited form as:

Stat Probab Lett. 2014 October 1; 93: 126–133. doi:10.1016/j.spl.2014.06.024.

A Modified Adaptive Lasso for Identifying Interactions in the Cox Model with the Heredity Constraint

Lu Wang^{a,*}, Jincheng Shen^a, and Peter F. Thall^b

^aDepartment of Biostatistics, University of Michigan, 1415 Washington Heights, Ann Arbor, MI 48109, USA

^bDepartment of Biostatistics, M.D. Anderson Cancer Center, Houston, Texas 77030, USA

Abstract

In many biomedical studies, identifying effects of covariate interactions on survival is a major goal. Important examples are treatment-subgroup interactions in clinical trials, and gene-gene or gene-environment interactions in genomic studies. A common problem when implementing a variable selection algorithm in such settings is the requirement that the model must satisfy the strong heredity constraint, wherein an interaction may be included in the model only if the interaction's component variables are included as main effects. We propose a modified Lasso method for the Cox regression model that adaptively selects important single covariates and pairwise interactions while enforcing the strong heredity constraint. The proposed method is based on a modified log partial likelihood including two adaptively weighted penalties, one for main effects and one for interactions. A two-dimensional tuning parameter for the penalties is determined by generalized cross-validation. Asymptotic properties are established, including consistency and rate of convergence, and it is shown that the proposed selection procedure has oracle properties, given proper choice of regularization parameters. Simulations illustrate that the proposed method performs reliably across a range of different scenarios.

Keywords

Modified adaptive Lasso; Oracle property; Penalized partial likelihood; Regularization; Variable selection

1. Introduction

The Cox proportional hazards regression model (Cox, 1972, 1975) is the most widely used statistical model for evaluating relationships between an event time, T , and baseline covariates, $X = (X_1, \dots, X_p)$. The Cox model is characterized by the hazard function

© 2014 Elsevier B.V. All rights reserved.

*Corresponding author. luwang@umich.edu (Lu Wang).

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

$$h(t|X, \beta) = h_0(t) \exp\{g(X, \beta)\}, \quad t > 0, \quad (1)$$

for a subject with covariates X , where $h_0(t)$ is an unspecified baseline hazard and the linear component $g(X, \beta) = \sum_{j=1}^p \beta_j X_j$ is characterized by a vector $\beta = (\beta_1, \dots, \beta_p)^T$ of unknown regression coefficients. In many applications, this simple form of $g(X, \beta)$ does not adequately describe the relationship between T and X due to interactions between elements of X . For example, an anti-cancer agent tailored to attack a certain biological target typically has effects hypothesized to differ between the subgroups of patients who do and do not have the target, identified by a binary “biomarker” covariate. In a randomised trial of the targeted agent versus standard therapy, such a differential effect is characterized as a treatment-biomarker interaction which, if found to be sufficiently large, may lead to regulatory approval of the agent for patients who are biomarker positive. Identifying such treatment-biomarker interactions thus is a key step in developing personalized treatments. The presence of other covariates that may or may not be associated with T , and that also may interact with treatment, complicates identification and estimation of biomarker effects. Fitting the Cox model with interactions is challenging since, even with a moderate number of covariates, the number of interaction terms may be large, and not all covariates and interactions may have meaningful effects on $h(t|X, \beta)$.

There is a large literature on variable selection methods for survival models. A family of penalized partial likelihood methods have been proposed for the Cox proportional hazard model, including the Lasso (Tibshirani et al., 1997) and the smoothly clipped absolute deviation method (Fan and Li, 2002). By shrinking some regression coefficients to zero, these methods simultaneously select important variables and estimate the regression model parameters. Zhang and Lu (2007) proposed an adaptive Lasso estimator for variable selection in the Cox model, with an adaptively weighted L_1 penalty on the regression coefficients that has a convex form. They showed that this method enjoys the oracle properties, global optima exist, and it can be implemented efficiently using standard numerical algorithms (Boyd and Vandenberghe, 2004).

All the above variable selection methods for the Cox model treat the candidate variables equally. However, when interactions are included, there is a natural hierarchical ordering among the variables in the model (Chipman, 1996; Joseph, 2006; Yuan et al., 2007). This motivates the strong heredity requirement that an interaction can be included in the model only if the interaction’s component variables are included as main effects (Hamada and Wu, 1992), since models that violate this property are difficult to interpret. For linear and generalized linear regression models, Yuan et al. (2009) proposed non-negative garrote methods that naturally incorporate a general hierarchical structure among predictors. Along this line, Choi et al. (2010) extended the Lasso to identify interaction terms while obeying the strong heredity constraint, which is achieved by reparameterizing the coefficients of the interaction terms. Bien et al. (2013) investigated a Lasso for hierarchical interactions, and Radchenko and James (2010) considered a more general case with nonlinear interactions. To our knowledge, none of these papers studied the setting with time-to-event data, and none of the variable selection methods for the Cox model noted above satisfy the strong heredity constraint when interactions are included. This paper aims at filling this gap.

We propose a modified Lasso procedure for the Cox model to adaptively select covariates and interactions while automatically enforcing the strong heredity constraint. The main challenges compared to linear/generalized linear regression are that the Cox model is semiparametric and involves right-censored data. We carry out estimation and variable selection by optimizing a modified log partial likelihood that includes two adaptively weighted penalty terms, one for main effects and one for reparameterized interactions, with each penalty multiplied by a tuning parameter. The main reason that we chose the reparameterization approach is that this method can handle all p main effects simultaneously in one iteration, which has tremendous numerical advantage, especially for survival analysis when no closed form can be found. The two-dimensional tuning parameter is determined by generalized cross-validation. The proposed method is computationally convenient, has good convergence properties, and implementation is straightforward. We establish asymptotic properties, including selection consistency, rate of convergence, and the oracle property (Fan and Li, 2001; Fan and Peng, 2004) that it performs as well as if the correct underlying model were known in advance. These theoretical properties and algorithms have not been studied previously for variable selection for Cox models with interaction terms subject to the strong hierarchy constraint.

2. Adaptive Lasso with Strong Heredity Constraint Using Penalized Partial Likelihood

Let T_i denote the failure time, C_i the censoring time, and $X_i = (X_{i1}, \dots, X_{ip})$ the covariate vector of the i^{th} subject, for $i = 1, \dots, n$, with $T_i = \min(T_i, C_i)$ and $\delta_i = I(T_i \leq C_i)$. Suppose that $(X_1, T_1, C_1), \dots, (X_n, T_n, C_n)$ are independent and identically distributed. We further assume non-informative censoring, $T_i \perp\!\!\!\perp C_i / X_i$, i.e., T_i is conditionally independent of C_i given X_i . Although our methods can be generalized to handle higher order interactions, for the ease of exposition we consider the Cox model having linear component with all possible two-way interactions. That is, in model (1)

$$g(X, \beta, \alpha) = \sum_{j=1}^p \beta_j X_j + \sum_{1 \leq j < j' \leq p} \alpha_{j,j'} X_j X_{j'}, \quad (2)$$

where $\alpha = (\alpha_{1,2}, \dots, \alpha_{p-1,p})^T$. Our goals are to provide a method that determines which terms in $g(X, \beta, \alpha)$ have important effects on the hazard, and develop a corresponding computational algorithm and parameter estimators having desirable properties. The existing variable selection methods for the Cox model do not guarantee the strong heredity constraint, as they treat all elements of (β, α) equally and do not distinguish between elements of β and α .

We first re-parameterize the coefficients for the interaction terms in (2) as $\alpha_{j,j'} = \gamma_{j,j'} \beta_j \beta_{j'}$, so that the linear term becomes

$$g(X, \theta) = \sum_{j=1}^p \beta_j X_j + \sum_{1 \leq j < j' \leq p} \gamma_{j,j'} \beta_j \beta_{j'} X_j X_{j'} \quad (3)$$

and the parameter vector is $\theta = (\beta, \gamma) = (\beta_1, \dots, \beta_p, \gamma_{1,2}, \dots, \gamma_{p-1,p})^T$. With this reparameterization, the coefficient for an interaction term $X_j X_{j'}$ must be 0 if either of its two main effects X_j or $X_{j'}$ has coefficient 0. Conversely, if $\gamma_{j,j'} \neq 0$, this implies that both $\beta_j \neq 0$ and $\beta_{j'} \neq 0$, which guarantees the strong heredity constraint. With the reparameterization (3), the log partial likelihood is

$$l_n(\theta) = \sum_{i=1}^n \delta_i \left(g(X_i, \theta) - \log \left[\sum_{r=1}^n I(\tilde{T}_r \geq \tilde{T}_i) \exp \{g(X_r, \theta)\} \right] \right), \quad (4)$$

where $I(A)$ denotes the indicator of the event A . For the variable selection problem at hand, we will minimize the adaptive penalized negative log partial likelihood

$$Q_n(\theta, \lambda_\beta, \lambda_\gamma) = -l_n(\theta) + n\lambda_\beta \sum_{j=1}^p w_j^\beta |\beta_j| + n\lambda_\gamma \sum_{1 \leq j < j' \leq p} w_{j,j'}^\gamma |\gamma_{j,j'}|, \quad (5)$$

where $\{w_j^\beta\}$ and $\{w_{j,j'}^\gamma\}$ are prespecified weights, and $\lambda_\beta, \lambda_\gamma$ are tuning parameters. Following Breiman (1995) and Zou (2006), we compute the weights in (5) using

$$w_j^\beta = \frac{\log(n)}{n} \left| \frac{1}{\tilde{\beta}_j} \right|, \quad w_{j,j'}^\gamma = \frac{\log(n)}{n} \left| \frac{\tilde{\beta}_j \tilde{\beta}_{j'}}{\tilde{\alpha}_{j,j'}} \right|,$$

where $\tilde{\beta}_j$'s and $\tilde{\alpha}_{j,j'}$'s are the estimates from a usual, unpenalized fitted Cox model with linear component (2). The multiplier $\log(n)/n$ is included to satisfy the convergence rate and the asymptotic properties introduced in Section 3, below. The objective function becomes

$$Q_n(\theta, \lambda_\beta, \lambda_\gamma) = -l_n(\theta) + \log(n) \lambda_\beta \sum_{j=1}^p \frac{|\beta_j|}{|\tilde{\beta}_j|} + \log(n) \lambda_\gamma \sum_{1 \leq j < j' \leq p} |\gamma_{j,j'}| \left| \frac{\tilde{\beta}_j \tilde{\beta}_{j'}}{\tilde{\alpha}_{j,j'}} \right|. \quad (6)$$

The two tuning parameters in (6) control the coefficient estimates at different levels. The first tuning parameter λ_β controls main effect estimates. If β_j is shrunk to zero, all terms involving X_j , including $\beta_j X_j$ and the interactions $\gamma_{j,j'} \beta_j \beta_{j'} X_j X_{j'}$, for any j' , are removed from the model. The second tuning parameter λ_γ controls the estimates only at the interaction effect level. Even if both $\beta_j \neq 0$ and $\beta_{j'} \neq 0$, it is possible that $\gamma_{j,j'} = 0$ if X_j and $X_{j'}$ do not interact. The penalty term controlled by λ_γ thus provides the flexibility of selecting only main effects of X_j and $X_{j'}$ but not their interaction.

The weights act on the objective function, $Q_n(\theta, \lambda_\beta, \lambda_\gamma)$, as follows. If the initial estimate $\tilde{\beta}_j$ is close to 0 then w_j^β will be large and hence, as can be seen from (5), the coefficient β_j of X_j will be heavily penalized. Similarly, if $\tilde{\alpha}_{j,j'}$ is small relative to $\tilde{\beta}_j \tilde{\beta}_{j'}$ then $w_{j,j'}^\gamma$ will be large and the coefficient $\gamma_{j,j'}$ of the interaction term $X_j X_{j'}$ will be heavily penalized.

3. Theoretical Properties of the Estimator

In this section, we study the asymptotic properties of our proposed variable selection procedure and the corresponding estimator. As $n \rightarrow \infty$, our estimator possesses the oracle property under certain regularity conditions, that is, it performs as well as if the true model were known in advance (Fan and Li, 2001). The regularity conditions that we need throughout the development are given in Appendix 1 in the web supplementary materials.

Let $\theta_0 = (\beta_0^T, \gamma_0^T)^T$ denote the true parameter vector, where

$$\gamma_{0j,j'} = \begin{cases} \alpha_{0j,j'} / \beta_{0j} \beta_{0j'} & \text{if } \beta_{0j} \neq 0 \text{ and } \beta_{0j'} \neq 0 \\ 0 & \text{otherwise} \end{cases} \quad (7)$$

This guarantees that the true model obeys the strong heredity constraint, that is, $\alpha_{0j,j'} = \gamma_{0j,j'} = 0$ if either $\beta_{0j} = 0$ or $\beta_{0j'} = 0$.

Define the covariate-specific tuning parameters $\lambda_{j,n}^\beta = \lambda_\beta w_j^\beta = n^{-1} \log(n) \lambda_\beta / |\tilde{\beta}_j|$ for $j = 1, \dots, p$ and the interaction-specific tuning parameters $\lambda_{j,j',n}^\gamma = \lambda_\gamma w_{j,j'}^\gamma = n^{-1} \log(n) \lambda_\gamma / |\tilde{\gamma}_{j,j'}|$ for $1 \leq j < j' \leq p$, where $\tilde{\gamma}_{j,j'} = \tilde{\alpha}_{j,j'} / \tilde{\beta}_j \tilde{\beta}_{j'}$. Our proposed estimator is

$$\hat{\theta} = \arg \min_{\theta} \left\{ -l_n(\theta) + n \sum_{j=1}^p \lambda_{j,n}^\beta |\beta_j| + n \sum_{1 \leq j < j' \leq p} \lambda_{j,j',n}^\gamma |\gamma_{j,j'}| \right\}. \quad (8)$$

Note that the criterion in (8) is the same as $Q_n(\theta, \lambda_\beta, \lambda_\gamma)$ in (6) with different notation.

Without loss of generality, we denote $\beta_0 = (\beta_{a0}^T, \beta_{b0}^T)^T$, where β_{a0} consists of all nonzero components and β_{b0} consists of the remaining zero components of β_0 . Similarly, for the true coefficient of interactions, write $\gamma_0 = (\gamma_{a0}^T, \gamma_{b0}^T, \gamma_{c0}^T)^T$, where γ_{a0} contains all nonzero components of γ_0 , γ_{b0} contains the zero components of γ_0 whose corresponding main effects both are nonzero, and γ_{c0} contains the remaining zero components of γ_0 whose corresponding components of main effects have at least one zero. Correspondingly, we

denote the maximizer of (8) as $\hat{\theta}_n = (\hat{\beta}_{an}^T, \hat{\beta}_{bn}^T, \hat{\gamma}_{an}^T, \hat{\gamma}_{bn}^T, \hat{\gamma}_{cn}^T)^T$, the covariate-specific tuning parameters as $\left\{ (\lambda_{an}^\beta)^T, (\lambda_{bn}^\beta)^T \right\}$, and the interaction-specific tuning parameters as $\left\{ (\lambda_{an}^\gamma)^T, (\lambda_{bn}^\gamma)^T, (\lambda_{cn}^\gamma)^T \right\}$.

Let

$$\xi_n = \max \left\{ (\lambda_{an}^\beta)^T, (\lambda_{an}^\gamma)^T \right\}, \zeta_n = \min \left\{ (\lambda_{bn}^\beta)^T, (\lambda_{bn}^\gamma)^T \right\}.$$

Then ξ_n is the maximum among both covariate-specific and interaction-specific tuning parameters that correspond to nonzero coefficients in the model. But for ζ_n , we only consider those specific tuning parameters corresponding to zero main effects and zero interaction terms when coefficients are in γ_{b0} . That is, we consider all zero cases for the main effects, but for the interactions, we only consider those zero terms whose corresponding main effects are nonzeros. We will refer to such terms in the definition of ζ_n as non-trivial zero terms.

For the proposed variable selection procedure to perform properly, as $n \rightarrow \infty$ the covariate-specific and interaction-specific penalties for the terms whose true coefficients are nonzeros should converge to 0, and the penalties for those non-trivial zero terms should be large enough so that the estimates shrink to 0. In fact, if the tuning parameters in (8) that correspond to β_{a0} and γ_{a0} converge to 0, and those corresponding to β_{b0} and γ_{b0} are sufficiently large, then our proposed estimating procedure will have the so-called oracle property (Fan and Li, 2001). This can be guaranteed when $\sqrt{n}$ -consistent estimates of θ_0 are used in the definition of $\lambda_{j,n}^\beta$ and $\lambda_{j,j',n}^\gamma$, where one can easily show that $\sqrt{n}\xi_n \rightarrow 0$ and $\sqrt{n}\zeta_n \rightarrow \infty$.

Denote the initial estimates as $\tilde{\beta} = (\tilde{\beta}_1, \dots, \tilde{\beta}_p)$, $\tilde{\gamma} = (\tilde{\gamma}_{1,2}, \dots, \tilde{\gamma}_{p-1,p})$, and we summarize the above results in Theorem 1. A sketch of the proof of Theorem 1 is given in Appendix 2.

Theorem 1. When $\sqrt{n}(\tilde{\beta} - \beta_0) = O_p(1)$ and $\sqrt{n}(\tilde{\gamma} - \gamma_0) = O_p(1)$ in (8) to calculate $\lambda_{j,n}^\beta$ and $\lambda_{j,j',n}^\gamma$, $\sqrt{n}\xi_n \rightarrow \infty$ and $\sqrt{n}\zeta_n \rightarrow \infty$ as $n \rightarrow \infty$. Under the regularity conditions (1)-(3) in Appendix 1, there exists a local minimizer $\hat{\theta}_n$ of $Q_n(\theta)$ such that

- i. (Sparsity) $P(\hat{\beta}_{bn}=0) \rightarrow 1$, $P(\hat{\gamma}_{bn}=0) \rightarrow 1$, and $P(\hat{\gamma}_{cn}=0) \rightarrow 1$ as $n \rightarrow \infty$.
- ii. (Asymptotic normality)

$$\sqrt{n} \left\{ \begin{pmatrix} \hat{\beta}_{an} \\ \hat{\gamma}_{an} \end{pmatrix} - \begin{pmatrix} \beta_{a0} \\ \gamma_{a0} \end{pmatrix} \right\} \xrightarrow{d} N \left\{ 0, \mathcal{J}_a^{-1}(\beta_{a0}, \gamma_{a0}) \right\},$$

where $\mathcal{J}_a(\beta_{a0}, \gamma_{a0})$ is the Fisher information matrix evaluated at β_{a0} and γ_{a0} assuming that $\beta_{b0} = 0$, $\gamma_{b0} = 0$, and $\gamma_{c0} = 0$ is known in advance.

Part (i) of Theorem 1 presents the sparsity property and shows that our proposed method can consistently remove the zero-effect terms with probability tending to 1. This implies that, with a sufficiently large sample, in practice our method can select the underlying true model with high probability. In part (ii) of Theorem 1, we establish that the estimates of nonzero elements of θ_0 are $\sqrt{n}$ -consistent and asymptotically normal. The asymptotic distribution is the same as what would be obtained if it were known in advance which elements of θ_0 are 0 and which are not 0, the so-called oracle property.

Theorem 2, below, establishes the asymptotic behavior of the proposed method when p is very large, especially when $p \rightarrow \infty$ as $n \rightarrow \infty$. This is an important case in many biomedical studies, with the advance of new high throughput technologies. A proof of Theorem 2 is given in the Appendix 3, provided in the web supplementary materials.

When the number of predictors may increase with the sample size n , we denote p as p_n , which allows the possibility that $p_n \rightarrow \infty$ as $n \rightarrow \infty$. Including all main effects and pairwise interactions, the total number of parameters is $q_n = (p_n + 1)p_n/2$. Similarly, when appropriate we add a subscript n to other notation, and we let d_n denote the number of non-zero coefficients in the underlying true model. Then we have

Theorem 2. Under the regularity conditions (4)-(6) in Appendix 1, if $p_n = o(n^{1/5})$ and

$\sqrt{nqn}\xi_n \rightarrow 0$, $\sqrt{n/qn}\zeta_n \rightarrow \infty$ as $n \rightarrow \infty$, then there exists a local minimizer $\hat{\theta}_n$ of $Q_n(\theta)$ such that

- i. (Sparsity) $P(\hat{\beta}_{bn}=0) \rightarrow 1$, $P(\hat{\gamma}_{bn}=0) \rightarrow 1$, and $P(\hat{\gamma}_{cn}=0) \rightarrow 1$ as $n \rightarrow \infty$.
- ii. (Asymptotic Normality)

$$\sqrt{n}\Omega_n\mathcal{I}_{an}^{1/2}(\beta_{a0}, \gamma_{a0}) \left\{ \begin{pmatrix} \hat{\beta}_{an} \\ \hat{\gamma}_{an} \end{pmatrix} - \begin{pmatrix} \beta_{a0} \\ \gamma_{a0} \end{pmatrix} \right\} \xrightarrow{d} N\{0, \Sigma\},$$

where Ω_n is any arbitrary $d \times d_n$ matrix with a finite d such that $\Omega_n\Omega_n^T \rightarrow \Sigma$, Σ is a $d \times d$ semipositive definite symmetric matrix, and $\mathcal{I}_{an}(\beta_{a0}, \gamma_{a0})$ is the $d_n \times d_n$ Fisher information matrix evaluated at $(\beta_{a0}, \gamma_{a0})$ assuming that $\beta_{b0} = 0$, $\gamma_{b0} = 0$, and $\gamma_{c0} = 0$ is known in advance.

The reason we consider an arbitrary linear combination $\Omega_n(\hat{\beta}_{an}^T, \hat{\gamma}_{an}^T)^T$ in Theorem 2, instead of $(\hat{\beta}_{an}^T, \hat{\gamma}_{an}^T)^T$ as in Theorem 1, is because the latter has dimension $d_n \rightarrow \infty$ as $n \rightarrow \infty$ in the current setup, while the former has finite dimension d .

4. Computational Algorithm

Expanding $l_n(\theta)$ in the objective function (6), $Q_n(\theta, \lambda_\beta, \lambda_\gamma)$ becomes

$$\begin{aligned} & -\sum_{i=1}^n \delta_i \left[\left(\sum_{j=1}^p \beta_j X_{ij} + \sum_{1 \leq j < j' \leq p} \gamma_{j,j'} \beta_j \beta_{j'} X_{ij} X_{ij'} \right) - \log \left\{ \sum_{r=1}^n I(\tilde{T}_r \geq \tilde{T}_i) \times \exp \left(\sum_{j=1}^p \beta_j X_{rj} + \sum_{1 \leq j < j' \leq p} \gamma_{j,j'} \beta_j \beta_{j'} X_{rj} X_{rj'} \right) \right\} \right] \\ & + \log(n) \lambda_\beta \sum_{j=1}^p \frac{|\beta_j|}{|\tilde{\beta}_j|} + \log(n) \lambda_\gamma \sum_{1 \leq j < j' \leq p} |\gamma_{j,j'}| \left| \frac{\tilde{\beta}_j \tilde{\beta}_{j'}}{\tilde{\alpha}_{jj'}} \right| \end{aligned}$$

Our estimator minimizes $Q_n(\theta, \lambda_\beta, \lambda_\gamma)$, that is, $\hat{\theta}_n = \argmin_{\theta} Q_n(\theta, \lambda_\beta, \lambda_\gamma)$. To carry out the optimization, we apply a modified version of the iteratively reweighted least squares

algorithm (Green, 1984) with weighted L_1 penalties. Denoting the gradient vector of the partial likelihood by $\dot{l}(\theta) = -\partial l_n(\theta) / \partial \theta$ and the Hessian matrix by $\ddot{l}(\theta) = -\partial^2 l_n(\theta) / \partial \theta \partial \theta^T$, and using the Cholesky decomposition $\ddot{l}(\theta) = MM^T$, where M is an invertible lower triangular matrix, we define the pseudo response vector $Y = (M)^{-1} \{ \ddot{l}(\theta) \theta - \dot{l}(\theta) \}$. By the usual second-order Taylor expansion, $-l_n(\theta)$ in (6) can be approximated by the quadratic form $(Y - M^T \theta)^T (Y - M^T \theta) / 2$, and at each penalized iteratively reweighted least squares iteration we minimize

$$\frac{1}{2} (Y - M^T \theta)^T (Y - M^T \theta) + \log(n) \lambda_\beta \sum_{j=1}^p \frac{|\beta_j|}{|\tilde{\beta}_j|} + \log(n) \lambda_\gamma \sum_{1 \leq j < j' \leq p} |\gamma_{j,j'}| \left| \frac{\tilde{\beta}_j \tilde{\beta}_{j'}}{\tilde{\alpha}_{j,j'}} \right|. \quad (9)$$

Because the main effect coefficients β and the interaction coefficients γ are controlled at different levels, in each step we also iterate between these two sets, first fixing β to estimate γ , then fixing γ to estimate β , and iterating until convergence. This algorithm is guaranteed to converge, since the objective function decreases at each step. When β is fixed, the optimization in γ becomes a Lasso problem, hence one can use either the LARS/Lasso algorithm (Efron et al., 2004) or quadratic programming to solve for γ efficiently. When γ is fixed, we solve for $\beta_1, \dots, \beta_p$ sequentially. For each $j = 1, \dots, p$, we fix γ and $\beta_{[-j]} = (\beta_1, \dots, \beta_{j-1}, \beta_{j+1}, \dots, \beta_p)$, and the optimization becomes a simple Lasso problem with only one parameter, β_j . This is similar to the shooting algorithm (Fu, 1998; Zhang and Lu, 2007; Friedman et al., 2007). The tuning parameters are selected by minimizing the generalized cross validation statistic,

$$C_{pseudo-GCV} = \frac{l(\hat{\theta})}{(1 - df/n)^2},$$

over a reasonable range of λ_β and λ_γ , where df is the number of nonzero parameters in the fitted model.

For a fixed λ_β and λ_γ , the optimization algorithm proceeds as follows:

Step 1. Center and normalize each term $X_j, X_j X_{j'}, j < j'$, and $j, j' = 1, \dots, p$.

Step 2. Start with plausible initial values $\hat{\beta}_j^{(0)}$ and $\hat{\gamma}_{j,j'}^{(0)}, j < j'$, and $j, j' = 1, \dots, p$, such as the conventional Cox regression parameter estimates. Set $m=1$.

Step 3. Compute Y and M based on the current value $\hat{\theta}^{(m-1)}$. Denote

$$\hat{l}(\theta) = -\frac{1}{2} (Y - M^T \theta)^T (Y - M^T \theta).$$

Step 4. To update $\hat{\gamma}$, let $\hat{\gamma}^{(m)} = \arg \min_{\gamma} \sum_{i=1}^n \left\{ -\tilde{l} \left(\hat{\beta}^{(m-1)}, \gamma \right) + n \lambda_\gamma \sum_{j < j'} w_{j,j'}^\gamma |\gamma_{j,j'}| \right\}$.

Step 5. To update $\hat{\beta}$ for each $j = 1, \dots, p$ in sequence, let

$$\hat{\beta}_j^{(m)} = \arg \min_{\beta_j} \sum_{i=1}^n \left\{ -\tilde{l} \left(\hat{\beta}_{[-j]}^{(m-1)}, \beta_j, \hat{\gamma}^{(m)} \right) + \lambda_{\beta} w_j^{\beta} |\beta_j| \right\}.$$

Step 6. If the relative difference between $Q_n \left(\hat{\theta}^{(m-1)} \right)$ and $Q_n \left(\hat{\theta}^{(m)} \right)$,

$$\Delta^{(m)} = \frac{\left| Q_n \left(\hat{\theta}^{(m-1)} \right) - Q_n \left(\hat{\theta}^{(m)} \right) \right|}{\left| Q_n \left(\hat{\theta}^{(m-1)} \right) \right|},$$

is small enough, then stop. Otherwise, increment m to $m+1$ and return to Step 3.

Since $\gamma_{j,j'} = a_{j,j'} / \beta_j \beta_{j'}$, in Step 3 the minimization is actually over $\{a_{j,j'}\}$ with $\beta = \hat{\beta}^{(m-1)}$. This algorithm gives exact zeros for some coefficients and guarantees that the coefficients of the corresponding interactions are set to 0 whenever the corresponding is shrunk to 0. Based on our empirical experience, the algorithm converges quickly.

5. Numerical Results

In this section, we report results of a simulation study comparing our proposed method with the Lasso and adaptive Lasso, two popular variable selection methods for the Cox model, neither of which guarantees the strong heredity constraint. We follow the simulation setup in Zhang and Lu (2007), but also include two-way interaction terms as candidates for the variable selection. We consider sample size $n = 200$, and assume that there are $p = 5$ covariates of interest in each simulated dataset, denoted by $X_1, X_2, \dots, X_5$. Thus, the number of all possible two-way interaction terms is $p \times (p-1)/2 = 10$, and there are a total of 15 candidate terms. We assume each covariate follows a standard normal distribution, and consider two scenarios: (i) the covariates are independent, and (ii) the covariates have pairwise correlations $\text{Corr}(X_j, X_{j'}) = \rho^{|j-j'|}$, with $\rho = 0.5$. We generate the censoring times from a Uniform distribution having support $(0, \tau)$, with τ chosen to obtain a specified censoring rate of either 25% or 40%.

Then suppose failure times are generated from a Cox model with constant baseline hazard $\lambda_0 = 0.1$, where the coefficients of X_2, X_3 , and the interaction term $X_2 X_3$ are non-zero, and the other 12 coefficients are zero. We consider the following two models:

Model 1: $\beta_2^0 = -0.8, \beta_3^0 = -0.8$, and $\alpha_{2,3}^0 = -0.8$, corresponding to large effects.

Model 2: $\beta_2^0 = -0.3, \beta_3^0 = -0.3$, and $\alpha_{2,3}^0 = -0.3$, corresponding to small effects.

We simulate each case 100 times, and apply all three methods to each simulated data set. For all methods, the tuning parameters are selected using the generalized cross validation criterion $C_{\text{pseudo-GCV}}$.

Table 1 summarizes the percentage of fully correct variable selection among 100 replications for each method under each scenario. For all cases, the proposed method correctly selects the true model more frequently than the regular Lasso and adaptive Lasso. Specifically, for all scenarios under Model 2, where important variables and interactions have small effects, the Lasso and adaptive Lasso perform about the same, with the latter slightly better, while our proposed method performs much better than these two methods. In Model 1, where important variables and interactions have large effects, the advantage of our proposed method is still substantial when the censoring percentage is high, 40%. When censoring is moderate, 25%, the adaptive Lasso becomes competitive, but still is worse than the proposed method. The performance of the Lasso is always the worst among the three methods.

Table 2 summarizes the individual frequencies of being selected into the model for each main effect and interaction term. We only present the results here for 25% censoring under both models when the covariates are independent ($\rho = 0$) due to space limitations, but similar results are observed for the 40% censored case, and correlated covariates case. As shown in Table 2, the proposed variable selection procedure always chooses important variable and interaction terms much more often than the other two methods. Under Model 1 for large effects, the adaptive Lasso and the proposed method in general select the unimportant terms less often than the Lasso, while for Model 2 with small effects, the proposed method actually selects some of the unimportant terms more often. However, as illustrated in Table 1, the proposed method simultaneously selects all the important terms and removes the unimportant ones more frequently.

A simple, ad hoc alternative approach often employed in practice for variable selection involving interaction terms is to run either the Lasso or the adaptive Lasso, and then, depending on which terms were selected, manually add back any main terms that were not selected but that were components of any interaction term that was selected. This ensures the strong heredity constraint. As seen in our simulation, this practice does not help much in terms of how often the true model is selected correctly, while the false positive error rate is greatly increased. For example, as one can see in Table 2 for model 2, since the proposed method selected the true interaction term x_2x_3 more often than the other two methods, using the above ad hoc approach following the Lasso or the adaptive Lasso still cannot beat the proposed method. For model 1 with large effects in Table 2, the frequency of x_2 , x_3 , and x_2x_3 being correctly selected already achieves 100%, so this ad hoc approach does not help at all.

To measure the prediction accuracy, we average the mean squared error

$C_{MSE} = (\hat{\theta} - \theta)^T \Gamma (\hat{\theta} - \theta)$ over 100 replications by following Tibshirani et al. (1997) and Zhang and Lu (2007), where Γ is the population variance-covariance matrix of the covariates. Standard errors are given in parentheses. For all scenarios, the proposed method has the smallest mean squared error (Table 3), and thus outperform the other two competitors in terms of prediction accuracy.

6. Discussion

We have extended the adaptive Lasso method to accommodate the Cox proportional hazard model including interaction terms while ensuring that the strong heredity constraint is satisfied in the selected model. Hamada and Wu (1992) consider other constraints, such as weak heredity where only one of the two main terms is required to be included when an interaction term is considered. Although this situation is not of our main interest, our methods can easily be modified to handle it by employing a different reparameterization.

Similarly, as discussed in Zhang and Lu (2007), the adaptive choice of weights may become problematic when some elements of θ are not estimable. This may occur, for example, when strong collinearity exists among covariates, or when the number of covariates p is much larger than the sample size n in high-dimensional data. In such settings, one cannot obtain the initial estimates to determine the adaptive weights. Alternatively, robust estimation of θ , such as ridge regression, may be considered.

Our work was motivated by the desire to identify key treatment-biomarker interactions for developing personalized treatments, where the number of candidate biomarkers is usually fixed or may grow slowly with n . Even with a moderate number of candidate biomarkers, however, our methodology can have an important impact on physician's actual behavior, if clinically meaningful treatment-biomarker effects are identified. The condition $p = o(n^{1/5})$ that we used may be relaxed though. There have been recent theoretical developments on high-dimensional survival models such as those in Bradic et al. (2011) and Lin and Lv (2013), but their underlying theory cannot be applied directly to our setting, which uses a reparameterization approach for selecting important interactions with heredity constraint. In order to relax our condition $p = o(n^{1/5})$, modification of the above theory to our setting would be required, and this would entail a substantial amount of work. While this is beyond the scope of the current paper, it certainly would be an interesting future research project.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

The authors thank the Co-Editor-in-Chief, Associate Editor, and referee for their valuable comments.

References

- Andersen PK, Gill RD. Cox's regression model for counting processes: a large sample study. *The annals of statistics*. 1982;1100–1120.
- Bien J, Taylor J, Tibshirani R. A Lasso for Hierarchical Interactions. *The Annals of Statistics*. 2013; 41:1111–1141.
- Boyd, S.; Vandenberghe, L. *Convex Optimization*. Cambridge University Press; Cambridge: 2004.
- Bradic J, Fan J, Jiang J. Regularization for Cox's Proportional Hazards Model With NP-Dimensionality. *The Annals of Statistics*. 2011; 39(6):3092–3120.
- Breiman L. Better Subset Regression Using the Non-Negative Garrote. *Technometrics*. 1995; 37:373–384.

- Chipman H. Bayesian variable selection with related predictors. *Canadian Journal of Statistics*. 1996; 24:17–36.
- Choi NH, Li W, Zhu J. Variable selection with the strong heredity constraint and its oracle property. *Journal of the American Statistical Association*. 2010; 105:354–364.
- Cox DR. Regression models and life tables (with discussion). *Journal of the Royal Statistical Society*. 1972; 34:187–220.
- Cox DR. Partial likelihood. *Biometrika*. 1972; 1975; 62:269–276.
- Efron B, Hastie T, Johnstone I, Tibshirani R. Least angle regression (with Discussion). *The Annals of statistics*. 2004; 32:407–499.
- Fan J, Li R. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association*. 2001; 96:1348–1360.
- Fan J, Li R. Variable selection for Cox’s proportional hazards model and frailty model. *The Annals of Statistics*. 2002; 30:74–99.
- Fan J, Peng H. Nonconcave penalized likelihood with a diverging number of parameters. *The Annals of Statistics*. 2004; 32:928–961.
- Friedman J, Hastie T, Höfling H, Tibshirani R. Pathwise coordinate optimization. *The Annals of Applied Statistics*. 2007; 1:302–332.
- Fu WJ. Penalized regressions: the bridge versus the lasso. *Journal of computational and graphical statistics*. 1998; 7:397–416.
- Green PJ. Iteratively reweighted least squares for maximum likelihood estimation, and some robust and resistant alternatives. *Journal of the Royal Statistical Society. Series B (Methodological)*. 1984:149–192.
- Hamada M, Wu CJ. Analysis of designed experiments with complex aliasing. *Journal of Quality Technology*. 1992; 24:130–137.
- Joseph VR. A Bayesian approach to the design and analysis of fractionated experiments. *Technometrics*. 2006; 48:219–229.
- Lin W, Lv J. High-Dimensional Sparse Additive Hazards Regression. *Journal of the American Statistical Association*. 2013; 108:247–264.
- McCullagh, P.; Nelder, JA. Generalized linear model. Vol. 37. Chapman & Hall/CRC; 1989.
- Nelder JA. The statistics of linear models: back to basics. *Statistics and Computing*. 1994; 4:221–234.
- Radchenko P, James GM. Variable Selection Using Adaptive Nonlinear Interaction Structures in High Dimensions. *Journal of the American Statistical Association*. 2010; 105:1541–1553.
- Tibshirani R, et al. The lasso method for variable selection in the Cox model. *Statistics in medicine*. 1997; 16:385–395. [PubMed: 9044528]
- Yuan M, Joseph VR, Lin Y. An efficient variable selection approach for analyzing designed experiments. *Technometrics*. 2007; 49:430–439.
- Yuan M, Joseph VR, Zou H. Structured variable selection and estimation. *The Annals of Applied Statistics*. 2009:1738–1757.
- Zhang HH, Lu W. Adaptive Lasso for Cox’s proportional hazards model. *Biometrika*. 2007; 94:691–703.
- Zou H. The adaptive lasso and its oracle properties. *Journal of the American statistical association*. 2006; 101:1418–1429.

Table 1

Percentages of correctly selected models among 100 replications.

	Censoring Percentage	ρ for Correlation	Method		
			Lasso	Adaptive Lasso	Proposed Method
Model 1	25%	0	11%	70%	86%
		0.5	6%	66%	79%
	40%	0	3%	58%	80%
		0.5	3%	54%	82%
Model 2	25%	0	43%	49%	74%
		0.5	26%	31%	63%
	40%	0	28%	32%	74%
		0.5	22%	33%	62%

Table 2

Term-specific percentages of being selected for each main effect and interaction. (25% censoring and $\rho = 0$ for independent covariates)

	Method	x_1	x_2	x_3	x_4	x_5	$x_1 x_2$	$x_1 x_3$	$x_1 x_4$
Model 1	Lasso	22%	100%	100%	18%	25%	15%	12%	15%
	adaptive Lasso	6%	100%	100%	3%	7%	2%	4%	1%
	proposed method	5%	100%	100%	4%	6%	5%	5%	0%
Model 2	Lasso	6%	91%	91%	6%	5%	2%	6%	3%
	adaptive Lasso	2%	87%	85%	4%	4%	1%	2%	1%
	proposed method	10%	98%	99%	11%	8%	10%	10%	1%

	Method	$x_1 x_5$	$x_2 x_3$	$x_2 x_4$	$x_2 x_5$	$x_3 x_4$	$x_3 x_5$	$x_4 x_5$
Model 1	Lasso	15%	100%	13%	13%	11%	13%	15%
	adaptive Lasso	2%	100%	3%	2%	2%	2%	0%
	proposed method	1%	100%	4%	6%	4%	6%	0%
Model 2	Lasso	3%	89%	4%	3%	9%	5%	2%
	adaptive Lasso	2%	82%	1%	2%	1%	0%	0%
	proposed method	1%	97%	10%	8%	11%	8%	2%

Table 3

Mean squared error, with standard errors in parentheses.

	Censoring Percentage	ρ for Correlation	Method		
			Lasso	Adaptive Lasso	Proposed Method
Model 1	25%	0	0.213 (0.124)	0.072 (0.068)	0.063 (0.065)
		0.5	0.219 (0.136)	0.071 (0.068)	0.069 (0.065)
	40%	0	0.337 (0.209)	0.168 (0.144)	0.117 (0.219)
		0.5	0.351 (0.215)	0.162 (0.172)	0.118 (0.216)
Model 2	25%	0	0.118 (0.053)	0.105 (0.055)	0.042 (0.041)
		0.5	0.131 (0.059)	0.130 (0.059)	0.054 (0.048)
	40%	0	0.117 (0.053)	0.122 (0.058)	0.061 (0.062)
		0.5	0.137 (0.066)	0.142 (0.068)	0.082 (0.078)