SEMIPARAMETRIC SINGLE-INDEX MODELS FOR OPTIMAL TREATMENT REGIMENS WITH CENSORED OUTCOMES

Jin Wang

A dissertation submitted to the faculty at the University of North Carolina at Chapel Hill in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the Department of Biostatistics in the Gillings School of Global Public Health.

Chapel Hill
2019

Approved by:

Danyu Lin

Donglin Zeng

Federico Innocenti

Michael Hudgens

Quefeng Li

ABSTRACT

Jin Wang: Semiparametric Single-Index Models for Optimal Treatment Regimens With Censored Outcomes
(Under the direction of Danyu Lin and Donglin Zeng)

There is a growing interest in precision medicine, where a potentially censored survival time is often the most important outcome of interest. To discover optimal treatment regimens for such an outcome, we propose a semiparametric proportional hazards model by incorporating the interaction between treatment and a single index of covariates through an unknown monotone link function. This model is flexible enough to allow non-linear treatment-covariate interactions and yet provides a clinically interpretable linear rule for treatment decision. We propose a sieve maximum likelihood estimation approach, under which the baseline hazard function is estimated nonparametrically and the unknown link function is estimated via monotone quadratic B-splines. We show that the resulting estimators are consistent and asymptotically normal with a covariance matrix that attains the semiparametric efficiency bound. The optimal treatment rule follows naturally as a linear combination of the maximum likelihood estimators of the model parameters. Through extensive simulation studies and an application to an AIDS clinical trial, we demonstrate that the treatment rule derived from the single-index model outperforms the treatment rule under the standard Cox proportional hazards model.

We extend the proposed method to transformation models so that optimal treatment rules can be applied to flexible hazards relationships. The transformation model introduces new challenges to both the estimation procedure and the asymptotic properties of the estimators. We design an estimation procedure with the EM algorithm by recognizing the transformation function as the distribution function of a corresponding missing random variable. We prove

that the resulting estimators are consistent and asymptotically normal, with the covariance matrix estimated using the profile likelihood theory. We demonstrate the performance of the transformation single-index model in simulation studies. We show that the proposed treatment rule under the single-index transformation model is more effective than that under the single-index proportional hazards model in delaying the disease relapse of large-bowel carcinoma in a real data analysis.

With improvements in technology, researchers are able to collect many clinical and genetic variables; not all the covariates may contribute to the prediction of the optimal treatment rules. We apply the adaptive Lasso penalty to the log-likelihood of the proposed model and let the data automatically determine the important predictors in the optimal treatment regime. We propose a simple computational approach by quadratic approximation of the original objective function and utilization of the variable selection software package available for the proportional hazards model. We show that the proposed variable selection approach displays the oracle property. The performance of the variable selection procedure is demonstrated in extensive simulations and the analysis of a multi-cancer clinical trial.

To my parents

ACKNOWLEDGEMENTS

My deepest gratitude is to my advisors, Dr. Danyu Lin and Dr. Donglin Zeng. My research could not have been accomplished without their patient guidance, mentoring and support throughout the process.

I owe many thanks to the rest of my committee members for their insightful suggestions and helpful discussions at different stages of my research. I am thankful to Dr. Michael Hudgens for his comments on my dissertation and the kind encouragement. The collaboration experience with Dr. Federico Innocenti has been an amazing journey. My thanks also goes to Dr. Quefeng Li for his invaluable advice about my research.

Finally, none of this would be possible without the love and support from my parents, Fusheng Wang and Hanhong Jin. This dissertation is dedicated to them.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1: INTRODUCTION

There is tremendous variation in the way patients respond to medications, in terms of both toxicity and treatment efficacy. Thus, when a treatment has an average positive effect in a population, its benefit may vary across patients according to certain baseline characteristics. Such phenomenon is demonstrated in many studies.

For example, the International Breast Cancer Study Group launched a randomized trial to evaluate the role of adjuvant chemotherapy preceding treatment with tamoxifen for patients with lymph node-negative disease (Castiglione-Gertsch et al., 2002). This study showed that, in terms of the disease–free survival time, postmenopausal patients with lymph node-negative breast cancer benefited significantly from adjuvant chemotherapy if their cancer type was endocrine–nonresponsive. If their cancer was endocrine–responsive, they obtained no benefit from the combination treatment compared with tamoxifen alone.

In a more extreme situation, a treatment may be beneficial to some patients but harmful to others (Neumann et al., 2001). Researchers aimed to investigate whether or not treatment of C pneumoniae infection with antibiotics prevented restenosis after coronary stent placement. The findings showed that roxithromycin reduced the rate of restenosis after coronary stenting in patients with high C pneumonia titres, but it performed worse than placebo for patients with negative C pneumoniae titres. In other words, non-selective use of roxithromycin is inadequate for prevention of restenosis after coronary stenting.

Instead of administering the same treatment regimen to all patients with a particular disease, precision medicine aims to maximize clinical benefits by steering patients to the right drug at the right dose at the right time according to their clinical and genetic profiles (Lavori and Dawson, 2000; Hamburg and Collins, 2010). In various applications in different fields, a treatment can represent not only a drug or a dose, but also a policy, an intervention,

or a strategy. For example, economists are interested in predicting the outcome of each labor market program and suggesting an optimal labor market for each unemployeed job seeker (Behncke et al., 2009). Political science researchers are interested in the estimation of heterogeneous effects of different voter mobilization strategies (Imai et al., 2013). In this work, we design a novel approach to flexibly quantify the treatment-covariate interaction and guide patients to their optimal treatment.

CHAPTER 2: LITERATURE REVIEW

In this chapter, we review all the literatures related to our topic, and outline our proposed work.

In section 2.2, we review the existing works on estimation of the optimal treatment strategy with non-censored and censored outcomes. In section 2.3, we review the single index model and its application to data with censored outcomes. We conclude this chapter by outlining the proposed work in section 2.4.

2.1 Statistical Methods for Censored outcomes

There have been extensive studies on estimating the optimal treatment regimes for non-censored outcomes. There are three main types of approaches: Q-learning (Watkins, 1989; Watkins and Dayan, 1992; Murphy, 2005; Song et al., 2015), A-learning (Murphy, 2003; Robins, 2004), and model-free or policy search methods (Robins et al., 2008; Orellana et al., 2010).

Q-learning and A-learning estimate optimal dynamic treatment rules based on regression-type modelling. Q-learning involves postulating a regression model of the outcome of interest on treatment assignment and patient covariates, with the optimization steps implemented through a backward recursive fitting procedure. A-learning involves the same recursive strategy, but requires only a model on the contrasts among outcome with different treatments and the propensity scores, which are the probabilities of observed treatment assignment conditional on the patient information at each decision point. This may make A-learning more robust to model misspecification than Q-learning. Compared to Q-learning, which could be sensitive to model misspecification, A-learning is doubly robust in that the corresponding estimating equations are asymptotically unbiased as long as either the propensity score and

the outcome model is correctly specified.

A concern with both Q-learning and A-learning is the potential model misspecification on the regression relationship between the outcome and predictors. Zhao et al. (2009) proposed to use nonparametric regression techniques for the estimation of optimal rules. However, the resulting optimal rules may be complicated functions of patient information, which could be potentially high dimensional. As a result, the estimated treatment regime is difficult to interpret.

Alternatively, researchers have proposed the policy search methods. This class of methods directly derives and maximizes a consistent estimator for the value function over a restricted class of treatment regimes indexed by a finite number of parameters. Zhang et al. (2012) used the policy search approach to estimate the optimal regime within a prespecified class of treatment regimes by directly maximizing a doubly robust augmented inverse probability weighted estimator for the population summary of outcome over all regimes in the class. Zhang et al. (2013) adapted this approach to two or more decision point. The optimization step in such methods is usually challenging because the value estimator is non-smooth, which needs to be solved by non-standard optimization techniques. Zhao et al. (2012) and Zhang et al. (2012) recast this problem into the weighted classification framework and used readily-available classfication packages to solve the transformed problem. Matsouaka et al. (2014) employed a kernel smoothing technique to estimate non-parametrically the conditional mean for the difference of the potential outcomes in a subgroup of patients and derived its associated treatment regime.

Recently, there have been many other machine learning methods proposed to quantify the patient heterogeneity in response to treatment. For example, Wager and Athey (2018) extended the random forest method (Breiman, 2001) to model the heterogeneous treatment effects and estimate the optimal treatment regime. By building a large number of regression trees and averaging their predictions with data-driven weights, the random forest method allows for flexible modeling of potentially high dimensional interactions. Targeted Learning

(Van der Laan and Rose, 2011) was used in van der Laan and Luedtke (2015) and Luedtke and van der Laan (2016) to derive the two time-point treatment rule that maximizes the mean outcome of interest under the dynamic treatment regime. The authors proposed data adaptive estimators of the optimal treatment regime by sequential loss-based learning using both the blip function and weighted classification frameworks.

For all the machine learning methods discussed above, there are several limitations. First, black-box algorithms cannot completely provide interpretable understanding of the interactions between treatment and covariates. Second, there is lack of inference for assessing the distribution properties of the estimators. Risk bounds are usually crude for practical trials of small size. Third, the recently developed inference on estimators is not necessarily the most statistically efficient. Last, machine learning methods incorporate the censored outcomes by optimizing the truncated mean, optimizing survival probability at a prespecified time point, or discretizing the outcomes, which do not fully use all the information available. It is necessary for these methods to model the censoring distribution.

2.2 Statistical Methods for Personalized Medicine

2.2.1 Personalized Treatment Regime with Non-censored Outcomes

There have been extensive studies on estimating the optimal treatment regimes for non-censored outcomes. There are three main types of approaches: Q-learning (Watkins, 1989; Watkins and Dayan, 1992; Murphy, 2005; Song et al., 2015), A-learning (Murphy, 2003; Robins, 2004), and model-free or policy search methods (Robins et al., 2008; Orellana et al., 2010).

Q-learning and A-learning estimate optimal dynamic treatment rules based on regression-type modelling. Q-learning involves postulating a regression model of the outcome of interest on treatment assignment and patient covariates, with the optimization steps implemented through a backward recursive fitting procedure. A-learning involves the same recursive strategy, but requires only a model on the contrasts among outcome with different treatments

and the propensity scores, which are the probabilities of observed treatment assignment conditional on the patient information at each decision point. This may make A-learning more robust to model misspecification than Q-learning. Compared to Q-learning, which could be sensitive to model misspecification, A-learning is doubly robust in that the corresponding estimating equations are asymptotically unbiased as long as either the propensity score and the outcome model is correctly specified.

A concern with both Q-learning and A-learning is the potential model misspecification on the regression relationship between the outcome and predictors. Zhao et al. (2009) proposed to use nonparametric regression techniques for the estimation of optimal rules. However, the resulting optimal rules may be complicated functions of patient information, which could be potentially high dimensional. As a result, the estimated treatment regime is difficult to interpret.

Alternatively, researchers have proposed the policy search methods. This class of methods directly derives and maximizes a consistent estimator for the value function over a restricted class of treatment regimes indexed by a finite number of parameters. Zhang et al. (2012) used the policy search approach to estimate the optimal regime within a prespecified class of treatment regimes by directly maximizing a doubly robust augmented inverse probability weighted estimator for the population summary of outcome over all regimes in the class. Zhang et al. (2013) adapted this approach to two or more decision point. The optimization step in such methods is usually challenging because the value estimator is non-smooth, which needs to be solved by non-standard optimization techniques. Zhao et al. (2012) and Zhang et al. (2012) recast this problem into the weighted classification framework and used readily-available classfication packages to solve the transformed problem. Matsouaka et al. (2014) employed a kernel smoothing technique to estimate non-parametrically the conditional mean for the difference of the potential outcomes in a subgroup of patients and derived its associated treatment regime.

Recently, there have been many other machine learning methods proposed to quantify

the patient heterogeneity in response to treatment. For example, Wager and Athey (2018) extended the random forest method (Breiman, 2001) to model the heterogeneous treatment effects and estimate the optimal treatment regime. By building a large number of regression trees and averaging their predictions with data-driven weights, the random forest method allows for flexible modeling of potentially high dimensional interactions. Targeted Learning (Van der Laan and Rose, 2011) was used in van der Laan and Luedtke (2015) and Luedtke and van der Laan (2016) to derive the two time-point treatment rule that maximizes the mean outcome of interest under the dynamic treatment regime. The authors proposed data adaptive estimators of the optimal treatment regime by sequential loss-based learning using both the blip function and weighted classification frameworks.

For all the machine learning methods discussed above, there are several limitations. First, black-box algorithms cannot completely provide interpretable understanding of the interactions between treatment and covariates. Second, there is lack of inference for assessing the distribution properties of the estimators. Risk bounds are usually crude for practical trials of small size. Third, the recently developed inference on estimators is not necessarily the most statistically efficient. Last, machine learning methods incorporate the censored outcomes by optimizing the truncated mean, optimizing survival probability at a prespecified time point, or discretizing the outcomes, which do not fully use all the information available. It is necessary for these methods to model the censoring distribution.

2.2.2 Personalized Treatment Regime with Failure Outcomes

Potentially censored survival times, such as times to disease occurrence and death, are clinically more important than non-censored outcomes, but more challenging to deal with statistically. The estimation of optimal treatment regimes dealing with potentially censored outcomes is relatively less developed. To illustrate the methods, we define a few notations. Let T denote the survival time, A the treatment indicator (with values 1 versus 0), and $\boldsymbol{X}$ a p-vector of baseline covariates. The optimal treatment regime could be constructed by estimating the treatment-covariate interaction in the Cox (1972) proportional hazards model.

In this framework, the hazard function of T conditional on A and $\boldsymbol{X}$

$$\lambda(t|A, \boldsymbol{X}) = \lambda(t) \exp\{\boldsymbol{\alpha}^{\mathrm{T}}\boldsymbol{X} + A\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{X}\}, \tag{2.1}$$

where $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ are unknown regression parameters, and $\lambda(\cdot)$ is an arbitrary baseline function. However, parametric forms of interactions are inflexible and likely to be misspecified due to the inherent complexity of the interactions between treatment and covariates. The resulting treatment rule suffers from the modeling bias.

Goldberg and Kosorok (2012) generalized the Q-learning (Zhao et al., 2009, 2011) method by modeling the completely observed survival time and adjusting for censoring by inverse-probability-of-censoring weighting (Robins et al., 1994). Suppose the censoring time $C \in [0, \tau]$ with probability 1. Let T be the maximal number of decision time-points for a given multistage time-dependent decision problem and $\overline{T}$ denote the random number of stages for an individual ($\overline{T} < T$). For each $t = 1, \cdots, T$, let R_{t-1} denote the length of the interval between decision time-points $t-1$ and t with $R_0 = 0$. The goal was to find a policy π, which is defined as a sequence of deterministic decision rules, $\pi_1, ..., \pi_T$, that maximizes the truncated-by-τ mean survival time

$$E_\pi \left[\max \left(\sum_{t=1}^{\overline{T}} R_t, \tau \right) \right].$$

They developed a methodology to solve the backward recursion in reinforcement learning when the number and timing of stages are flexible and derived the associated finite sample risk bounds on the generalization error of the learned treatment regime.

Zhao et al. (2015) extended the outcome-weighted learning approach (Zhao et al., 2012) to accommodate potentially censored outcomes. They showed that maximization of the mean survival time is equivalent to solving a weighted misclassification problem with weights involving both the observed outcome and the inverse probability of censoring. They further proposed a doubly robust version of outcome-weighted learning so that the obtained individualized treatment rule is consistent for the optimal rule if either the survival model or the

censoring model is correct. In this way, it avoids potential model misspecification and the numerical instability with high censoring rate in Goldberg and Kosorok (2012).

Adopting the outcome-weighted learning approach, Bai et al. (2017a) showed that optimization of the patient survival probability could also be cast into the weighted classification framework. They derived a doubly robust estimator for the value of a treatment regime, and estimated the optimal rule using a value search method. This approach is extended by Jiang et al. (2017) to more than one decision point. Jiang et al. (2017) proposed a value-search method to maximize the survival probability to derive an optimal treatment rule that maximizes the t-year survival probability, with potential multiple stages of decision making. In this work, the authors introduced the inverse propensity score weighted (IPSW) and augmented inverse propensity score weighted (AIPSW) Kaplan-Meier estimators of the t-year survival probability, where the AIPSW estimator is doubly robust, which is consistent if either the treatment assignment probability or the censoring distribution is modeled correctly. Kernel smoothing techniques were used to smooth the estimator of the value before the optimization step, which improved the finite sample performance of the treatment rule. In both works, the time point at which the survival probability is optimized needs to be specified before the analysis, which could be highly subjective and does not take into account the complete survival profile. Jiang et al. (2017) extended Jiang et al. (2017) to maximize a prespecified function of the survival function, including the truncated mean survival time and median survival time, $t-$year survival probability as special cases. However, these quantities are still a single summary measuring the survival performance and may be incomplete to describe the whole picture of survival profiles.

Diaz et al. (2018) proposed methods for constructing an ensemble of decision functions for the optimal rule with survival outcomes. These ensembles are linear combinations of estimators from a user-supplied library, with the linear combination coefficients chosen as the minimizer of the cross-validated risk. The authors proposed a doubly robust loss function so that the estimated rules will remain optimal if either the survival time or the

censoring distribution is modeled correctly. By the no-free-lunch theorems (Wolpert, 2002) for supervised learning, the generalization error of ensembles constructed from a library were shown to be smaller or at least equal to the error of any individual candidate in the library. With a list of prespecified time points $\{1, \cdots, K\}$, the time-to-event outcome T takes values in $\{1, \cdots, K\} \cup \infty$, where $T = \infty$ denotes no event observed during time $[0, K]$. The censoring time C takes values in $\{1, \cdots, K\}$. Due to the assumed longitudinal data structure, the exact observation time is not used. Instead, information is only available on the patient status at the prespecified time points. Although there are certain studies that use such follow-up approach, in most of the clinical studies, the exact observed time is recorded. Throwing such information away results in information loss and the obtained estimator is not the most efficient one.

Kang et al. (2018) adapted an A-learning approach to estimate the optimal treatment regime estimation for censored data using a flexible additive hazards regression model. The additive hazard model allowed flexible baseline covariate effects in the control group and gave a closed form estimator for the optimal treatment regime. The authors modified the standard A-learning estimating equation (Robins et al., 1994) by introducing a time-dependent propensity score, which is defined as the probability of receiving the treatment for patients at risk conditional on their covariates. With the time-dependent propensity score, the estimator was shown to have improved robustness against misspecification of the baseline covariate effect model. A resampling method was proposed to estimate the asymptotic variance of the estimator. Extensions of this method to the Cox proportional hazards model is possible. However, the corresponding estimation is much more complicated and was not studied in this paper.

2.3 Single Index Models

2.3.1 Model and Estimation

The single index model was first proposed by Brillinger (1982) for flexible modeling of the outcome on covariates in the linear regression setting. Let Y denote the outcome of interest, $\boldsymbol{X}$ denote the covariates, and $\boldsymbol{\beta}$ denote the unknown coefficient. The ordinary least square regression problem takes the form

$$Y = \boldsymbol{X}^{\mathrm{T}}\boldsymbol{\beta} + \epsilon,$$

where $E(\epsilon|\boldsymbol{X}) = 0$ and ϵ is independent and identically distributed for each observation. Single index models were introduced to generalize linear regression by replacing the linear predictor with a semi-parametric component:

$$Y = \psi(\boldsymbol{X}^{\mathrm{T}}\boldsymbol{\beta}) + \epsilon,$$

where $E(\epsilon|\boldsymbol{X}) = 0$ and ψ is an arbitrary smooth function. For identifiability reasons, usually it is required that $||\boldsymbol{\beta}|| = 1$ and the first element of $\boldsymbol{\beta}$ is positive. Due to their flexibility and interpretability of the coefficients, as well as the ability to model flexible interactions among predictors, single index models have become increasingly popular in many scientific fields. Compared to alternative fully-nonparmetric methods such as the additive model, single index model effectively circumvents the curse of dimensionality.

There has been extensive research on the single model on the estimation procedure (Stoker, 1986; Powell et al., 1989; Duan and Li, 1991; Ichimura, 1993; Hardle et al., 1993; Weisberg and Welsh, 1994), and variable selection (Peng and Huang, 2011; Radchenko, 2015). The single index models were extended to generalized linear model framework (Carroll et al., 1997; Chiou and Müller, 1998; Chiou et al., 1999) and partially linear single-index models (Luo and Ghosal, 2016). However, extending the single index model to accommodate censored

outcomes has been challenging both theoretically and computationally. Popular models that deal with censored outcomes, such as the Cox proportional hazards model, contain a non-parametric baseline hazard function itself. When introducing the smooth link function into the model, the extra semiparametric component imposes challenges for the development of asymptotic theory. Estimation techniques in the standard single index model such as the slice inversion regression (Duan and Li, 1991) and the average derivative estimation (Stoker, 1986; Powell et al., 1989; Hardle et al., 1993) could not be directly applied. Therefore, novel algorithms is needed to compute the estimators.

2.3.2 Application of Single Index Models to Failure Outcomes

Single-index functions in the proportional hazards model have previously been used to account for nonlinear main effects of predictors on failure time outcomes. Wang (2004) proposed to use the single index function to relax the log-linear assumption in the Cox model. Let T and C be defined as in section 2.2.2. Let $\tilde{T} = \min\{T, C\}$ denote the observed time, $\Delta = I(T \leqslant C)$ denote the censoring indicator, and $\boldsymbol{Z}(t)$ denote the covariates which could be time-dependent. With the independent censoring assumption, the hazard function for the failure time $\widetilde{T}$ with covariates $\boldsymbol{Z}(t)$ under the Cox model is given by $\lambda(t|Z(t)) = \lambda_0(t)\exp\{\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z}(t)\}$. Wang (2004) considers a more general class of models with the hazard relationship

$$\lambda(t|Z(t)) = \lambda_0(t)\psi\{\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z}(t)\}$$

with $||\boldsymbol{\beta}|| = 1$, where ψ dose not necessarily take the log-linear form. When the covariate is one-dimensional time-independent, the form of the link function could be checked using alternative approaches in Tibshirani and Hastie (1987) and Fan et al. (1997). The methods by Wang (2004) could deal with time-dependent multi-covariates. The estimation for ψ followed from the local likelihood approach in Fan et al. (1997) by the p-th order Taylor expansion of $\psi\{\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z}(t)\}$ around $\psi\{\widetilde{\boldsymbol{\beta}}^{\mathrm{T}}\boldsymbol{Z}(t)\}$, where $\widetilde{\boldsymbol{\beta}}$ is an initial estimate and p is a prespecified order of expansion. With the estimated derivatives of ψ, $\widehat{\psi}$ is obtained by the trapezoidal rule

in numerical integration. Finally, $\boldsymbol{\beta}$ is estimated by optimizing the likelihood with the link function replaced with $\widehat{\psi}$.

Researchers have proposed various methods to model the (partially) nonparametric covariate effects on the log hazard ratio with potentially censored outcomes (Sasieni, 1992b,a; Nielsen et al., 1995; Huang et al., 1999; Lu et al., 2001). These methods suffer from the curse of dimentionality as the covariate effect is modeled in an unstructured nonparametric manner. Because of the sparsity of the data in even moderately large dimensions, accurate estimation under these models is usually difficult with the sample sizes in practical data set. Lu et al. (2006) proposed a class of partially linear single-index survival models in order to address this question with automatic dimension reduction achieved by the semiparametric single-index component. With the same notations in section 2.2.2, let covariates $\boldsymbol{X}$ be partitioned into two parts: p-dimensional covariates $\boldsymbol{X}_1$ assumed to have linear effects and q-dimensional covariates $\boldsymbol{X}_2$ assumed to have non-linear effects. Similarly, define $\boldsymbol{\beta}_1$ and $\boldsymbol{\beta}_2$ that correspond to $\boldsymbol{X}_1$ and $\boldsymbol{X}_2$. Let σ be an unknown parameter indexing the baseline hazard function. Lu et al. (2006) considered a class of semiparametric models defined by

$$\lambda(t; \boldsymbol{X}) = \lambda_0(t; \sigma) \exp\{\boldsymbol{\beta}_1^{\mathrm{T}} \boldsymbol{X}_1 + \psi(\boldsymbol{\beta}_2^{\mathrm{T}} \boldsymbol{X}_2)\},$$

with $||\boldsymbol{\beta}_2|| = 1$. Note that in this model, the baseline hazard function was modelled parametrically. The link function ψ is estimated using a quasi-likelihood approach using local linear kernels, which could be regarded as a generalization of Carroll et al. (1997). Due to the local kernel approximation approach, there is no shape constraint on the covariate effect. In addition, in this work as well as in Wang (2004), the authors treated the link function as a parametric component in the variance estimation and the asymptotic theory.

Huang and Liu (2006) considered an alternative way to model the possible nonlinearity of the covariate effects on the log hazard ratio in the proportional hazards model. With the

same notations defined in section 2.2.2, they specified the model as

$$\lambda(t|Z) = \lambda_0(t)\exp\{\psi(\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{X})\}$$

with $||\boldsymbol{\beta}|| = 1$. Unlike the local polynomial smoother approach adopted in the Wang (2004) and Lu et al. (2006), Huang and Liu (2006) estimated the unknown function ψ through a polynomial spline. With a chosen spline basis, the spline coefficients that corresponds to the basis expansion and the regression coefficients are simultaneously estimated by maximizing the partial likelihood. They treated the link function as a parametric component and derived asymptotic variance-covariance matrix following the partial likelihood theory in the standard Cox model. This approach is only valid conditional on the spline basis chosen in the estimation step. Since the fitted link function is not monotone, the interpretation of covariate effects is not straightforward.

2.4 Outline of Proposed Work

In the first project, we propose a flexible semiparametric single-index model to assess how the effect of treatment on the survival distribution depends on patient characteristics. The proposed model naturally extends the familiar Cox (1972) proportional hazards model by including the product of treatment and an arbitrary monotone function of the linear predictor. By using an arbitrary monotone regression function, the proposed model is flexible enough to allow complex treatment-covariate interactions while providing a computationally simple and clinically interpretable linear rule for personalized treatment decisions. To estimate the model parameters, we combine nonparametric maximum likelihood estimation with sieve estimation and develop an iterative alternating optimization procedure. We establish the asymptotic properties of the resulting estimators by novel applications of modern empirical process theory, seive estimation theory, and semiparametric efficiency theory. We demonstrate the usefulness of the proposed methods through simulation studies and an application to the

AIDS Clinical Trials Group (ACTG) 175 study (Hammer et al., 1996).

In the second project, we propose to extend the single-index model to transformation models to account for more flexible dependency of survival time on covariates. We show that if we cast the transformation function into the missing value framework, this problem could be solved by an EM algorithm. In each iteration, the maximization step is equivalent to optimizing a weighted proportional hazard model, with weights determined by posterior expectation of the missing data conditional on the observed data and the parameter value. We develop the associate asymptotic theory and show that the resulting estimators are semiparametrically efficient. The performance of the proposed methods is illustrated in simulation studies and a real data application to a clinical trial on patients with large-bowel carcinoma.

Last, we consider the variable selection issue in applying the single-index model to censored outcomes for the estimation of the optimal treatment regime. Recent technology advances make thousands or even millions of data available for each patient. In estimation of the optimal treatment rule, it is of great importance to tell the important feature from unimportant ones. Therefore the next question is to design a procedure that handles high-dimensional patient characteristics as input and select the relevant factors. We propose the adaptive lasso estimator for the single-index term to select such relevant coefficients, with the tuning parameter determined in a data-adaptive manner. We design an algorithm that computes the adaptive lasso estimate in a computationally efficient way. This proposed method is shown to have the oracle property. We conduct simulations to examine the performance of the variable selection procedure. A data application is included to show the effectiveness of the method.

CHAPTER 3: SINGLE-INDEX MODELS FOR OPTIMAL TREATMENT REGIMENS WITH THE COX MODEL

3.1 Introduction

There is tremendous variation in the way patients respond to medications, in terms of both toxicity and treatment efficacy. Thus, when a treatment has an average positive effect in a population, its benefit may vary across patients according to certain baseline characteristics. For example, the benefit of chemotherapy prior to hormone therapy with tamoxifen is much higher for postmenopausal women with lymph node–negative, estrogen receptor–negative breast cancer than for those with lymph node–negative, estrogen receptor–positive breast cancer (Castiglione-Gertsch et al., 2002). In a more extreme situation, a treatment may be beneficial to some patients but harmful to others (Neumann et al., 2001). Instead of administering the same treatment regimen to all patients with a particular disease, precision medicine aims to maximize clinical benefits by steering patients to the right drug at the right dose at the right time according to their clinical and genetic profiles (Lavori and Dawson, 2000; Hamburg and Collins, 2010).

Several methods are available to estimate personalized treatment strategies for non-censored outcomes, including parametric and semiparametric regression models (Murphy, 2003; Zhang et al., 2012, 2013) and machine learning methods (Zhao et al., 2012). Potentially censored survival times, such as times to disease occurrence and death, are clinically more important than non-censored outcomes, but more challenging to deal with statistically. The incorporation of treatment-covariate interactions into commonly used survival models, particularly the Cox (1972) proportional hazards model, can lead to estimation of the treatment strategy with the lowest hazard for each patient. However, parametric forms of interactions are inflexible and likely to be misspecified due to the inherent complexity of

the interactions between treatment and covariates. Several authors have modified machine learning methods to accommodate censoring. For example, Goldberg and Kosorok (2012) adopted the inverse-probability-of-censoring-weighting technique, which requires correct specification of the censoring distribution, and Zhao et al. (2015) derived outcome-weighted learning with doubly robust weights. These methods are focused on the (truncated) mean survival time, rather than the entire distribution function, and the treatment-covariate interactions are derived from a black box with no clinical interpretation.

Recently, Jiang et al. (2017) proposed a value-search method in order to derive an optimal treatment rule that maximizes the t-year survival probability. Their method requires estimation of the censoring distribution. By contrast, our approach does not require modeling the censoring distribution and provides the optimal treatment rules that lead to larger survival probabilities over time instead of at a pre-specified time point. In addition, our estimation procedure can be viewed as a value-search method for the optimal treatment rule even when the proportional hazards model is misspecified. In that case, the value function is the logarithm of the partial likelihood function, which essentially is the sum of the log-ratio between an exponential loss for each failure and the average loss for non-failures in the risk set. The parameters estimated by maximization of this value function distinguishes failures from non-failures the most and thus can lead to a beneficial treatment rule.

In this paper, we propose a flexible semiparametric single-index model to assess how the effect of treatment on the survival distribution depends on covariates. The proposed model naturally extends the familiar Cox (1972) proportional hazards model by including the product of treatment and an arbitrary monotone function of the linear predictor. By using an arbitrary monotone regression function, the proposed model is flexible enough to allow complex treatment-covariate interactions while providing a computationally simple and clinically interpretable linear rule for personalized treatment decisions. To estimate the model parameters, we combine nonparametric maximum likelihood estimation with sieve estimation and develop an iterative alternating optimization procedure. We establish the

asymptotic properties of the resulting estimators by novel applications of modern empirical process theory, seive estimation theory, and semiparametric efficiency theory. We demonstrate the usefulness of the proposed methods through simulation studies and an application to the AIDS Clinical Trials Group (ACTG) 175 study (Hammer et al., 1996).

3.2 Methods

3.2.1 Model and Optimal Treatment Regimens

Let T denote the survival time, A the treatment indicator (with values 1 versus 0), and $\boldsymbol{X}$ a p-vector of bounded baseline covariates. To allow the treatment effect to depend on covariates in a flexible manner, we propose a semiparametric single-index model in the proportional hazards form, such that the hazard function of T conditional on A and $\boldsymbol{X}$ is

$$\lambda(t|A, \boldsymbol{X}) = \lambda(t) \exp\{\boldsymbol{\alpha}^{\mathrm{T}}\boldsymbol{X} + A\psi(\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z})\}, \tag{3.1}$$

where $\boldsymbol{Z}$ is a q-dimensional subset of $\boldsymbol{X}$, $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ are unknown regression parameters, $\lambda(\cdot)$ is an arbitrary baseline function, and $\psi(\cdot)$ is an unknown and strictly increasing link function. For identifiability, we assume that there are non-zero components in $\boldsymbol{\beta}$ and $||\boldsymbol{\beta}|| = 1$, where $||\boldsymbol{x}||$ is the Euclidean norm of vector $\boldsymbol{x}$. If $\psi(u) = u$, then model 3.1 reduces to the standard Cox (1972) proportional hazards model with an interaction between treatment and covariates.

To see how the parameters in model (3.1) can yield the optimal treatment regimens, we let $T(a)$ denote the potential survival time if the patient receives treatment a for $a = 0$ or 1. We further make the stable unit treatment assumption and no unmeasured confounder assumption (Rubin, 1974):

(A1) $T = T(a)$ when $A = a$, and

(A2) A is independent of $\{T(0), T(1)\}$ conditional on $\boldsymbol{X}$.

Under these two assumptions, model (3.1) implies that

$$P(T(a) > t|\boldsymbol{X}) = P(T(a) > t|A = a, \boldsymbol{X}) = P(T > t|A = a, \boldsymbol{X})$$

$$= \exp\left\{-\Lambda(t)e^{\boldsymbol{\alpha}^{\mathrm{T}}\boldsymbol{X}+a\psi(\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z})}\right\}.$$

Thus, the survival probability of the potential survival time under $a^* \equiv I\{\psi(\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z}) < 0\}$, where $I(\cdot)$ is the indicator function, is always larger than the one under $(1 - a^*)$. In other words, the optimal treatment regimen should be $I(\psi(\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z}) < 0)$ or, equivalently, $I\{-\psi^{-1}(0) + \boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z} < 0\}$. Hence, the semiparametric single-index model provides a simple linear rule that can be conveniently used in practice.

3.2.2 Sieve Maximum Likelihood Estimation

Let C denote the censoring time, and write $\widetilde{T} = \min(T, C)$ and $\Delta = I(T \leqslant C)$. We assume that C is independent of T conditional on $(A, \boldsymbol{X})$. For a randomized clinical trial with n patients, the data consist of $(\widetilde{T}_i, \Delta_i, A_i, \boldsymbol{X}_i)$ $(i = 1, \ldots, n)$. The log-likelihood concerning the model parameters is given by

$$\sum_{i=1}^{n}\left[\Delta_i\left\{\boldsymbol{\alpha}^{\mathrm{T}}\boldsymbol{X}_i + A_i\psi(\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z}_i)\right\} + \Delta_i \log\lambda(\widetilde{T}_i) - \exp\{\boldsymbol{\alpha}^{\mathrm{T}}\boldsymbol{X}_i + A_i\psi(\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z}_i)\}\Lambda(\widetilde{T}_i)\right],$$

where $\Lambda(t) = \int_0^t \lambda(s)ds$.

We propose a sieve maximum likelihood estimation approach. First, we estimate the cumulative baseline hazard function $\Lambda(\cdot)$ nonparametrically by treating it as a step function with jumps at the observed survival times. Maximization of the above log-likelihood with respect to the jump sizes yields the profile log-likelihood

$$\sum_{i=1}^{n}\Delta_i\left(\boldsymbol{\alpha}^{\mathrm{T}}\boldsymbol{X}_i + A_i\psi(\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z}_i) - \log\left[\sum_{j=1}^{n}Y_j(\widetilde{T}_i)\exp\{\boldsymbol{\alpha}^{\mathrm{T}}\boldsymbol{X}_j + A_j\psi(\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z}_j)\}\right]\right), \tag{3.2}$$

where $Y_i(t) = I(\widetilde{T}_i \geqslant t)$.

We approximate $\psi(u)$ by B-splines (Schumaker, 1981). Specifically, let $\boldsymbol{B}(u) \equiv (B_1(u), \cdots, B_{K_n+1}(u))^{\mathrm{T}}$ denote quadratic B-spline bases corresponding to K_n distinct knots in an interval containing the union of the support of $\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z}$ for any unit vector $\boldsymbol{\beta}$. We then approximate $\psi(u)$ by $\widetilde{\psi}(u) \equiv \boldsymbol{\gamma}^{\mathrm{T}}\boldsymbol{B}(u)$, where $\boldsymbol{\gamma} \equiv (\gamma_1, \cdots, \gamma_{K_n+1})^{\mathrm{T}}$ is a vector of unknown

coefficients. To ensure that $\widetilde{\psi}(u)$ is increasing but not too large, we require that $-M_n \leqslant \gamma_1 \leqslant \cdots \leqslant \gamma_{K_n+1} \leqslant M_n$ for some constant M_n. The increasing sequence of γ values guarantees the increasing property of $\widetilde{\psi}(u)$ because of the choice of the quadratic splines.

Finally, to estimate $\boldsymbol{\alpha}, \boldsymbol{\beta}$, and $\boldsymbol{\gamma}$, we maximize

$$pl(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}) \equiv \sum_{i=1}^{n} \Delta_i \left(\boldsymbol{\alpha}^{\mathrm{T}} \boldsymbol{X}_i + A_i \boldsymbol{\gamma}^{\mathrm{T}} \boldsymbol{B}(\boldsymbol{\beta}^{\mathrm{T}} \boldsymbol{Z}_i) - \log \Big[\sum_{j=1}^{n} Y_j(\widetilde{T}_i) \exp\{\boldsymbol{\alpha}^{\mathrm{T}} \boldsymbol{X}_j + A_j \boldsymbol{\gamma}^{\mathrm{T}} \boldsymbol{B}(\boldsymbol{\beta}^{\mathrm{T}} \boldsymbol{Z}_j)\} \Big] \right) \quad (3.3)$$

under the constraints that $||\boldsymbol{\beta}|| = 1$ and $-M_n \leqslant \gamma_1 \leqslant \cdots \leqslant \gamma_{K_n+1} \leqslant M_n$. In the next section, we will describe the optimization algorithm and the choices of K_n, M_n, and the knots for defining B-splines.

3.2.3 Numerical Algorithm

We describe an iterative alternating optimization procedure (Bezdek and Hathaway, 2003) to maximize 3.3 as follows.

In the initial step, we fit the standard Cox model with covariates $(\boldsymbol{X}^{\mathrm{T}}, A, A\boldsymbol{Z}^{\mathrm{T}})^{\mathrm{T}}$, and we obtain initial estimates $\widehat{\boldsymbol{\alpha}}_0$ and $\widehat{\boldsymbol{\beta}}_0$ by using the estimated coefficients for $\boldsymbol{X}$ and the normalized coefficients for $A\boldsymbol{Z}$, respectively. The initial value for $\boldsymbol{\gamma}$ is set to the least-squares estimate that approximates a linear link using the B-splines, where the intercept and slope of this link function are, respectively, the coefficient of A and the Euclidean norm of the coefficients for $A\boldsymbol{Z}$ in the Cox model.

At the lth iteration, the current parameter values are denoted by $\widehat{\boldsymbol{\alpha}}_{l-1}, \widehat{\boldsymbol{\beta}}_{l-1}$, and $\widehat{\boldsymbol{\gamma}}_{l-1}$. We set the link function as $\widehat{\psi}_{l-1}(u) \equiv \widehat{\boldsymbol{\gamma}}_{l-1}^{\mathrm{T}} \boldsymbol{B}(u)$. We first update $\boldsymbol{\beta}$ by maximizing

$$\sum_{i=1}^{n} \Delta_i \Big\{ \widehat{\boldsymbol{\alpha}}_{l-1}^{\mathrm{T}} \boldsymbol{X}_i + A_i \{ \widehat{\psi}_{l-1}(\widehat{\boldsymbol{\beta}}_{l-1}^{\mathrm{T}} \boldsymbol{Z}_i) + \widehat{\psi}'_{l-1}(\widehat{\boldsymbol{\beta}}_{l-1}^{\mathrm{T}} \boldsymbol{Z}_i)(\boldsymbol{\beta} - \widehat{\boldsymbol{\beta}}_{l-1})^{\mathrm{T}} \boldsymbol{Z}_i \} - \log \Big(\sum_{j=1}^{n} Y_j(\widetilde{T}_i) \exp \Big[\widehat{\boldsymbol{\alpha}}_{l-1}^{\mathrm{T}} \boldsymbol{X}_j + A_j \{ \widehat{\psi}_{l-1}(\widehat{\boldsymbol{\beta}}_{l-1}^{\mathrm{T}} \boldsymbol{Z}_j) + \widehat{\psi}'_{l-1}(\widehat{\boldsymbol{\beta}}_{l-1}^{\mathrm{T}} \boldsymbol{Z}_j)(\boldsymbol{\beta} - \widehat{\boldsymbol{\beta}}_{l-1})^{\mathrm{T}} \boldsymbol{Z}_j \} \Big] \Big) \Big\}$$

subject to the constraint $\|\boldsymbol{\beta}\| = 1$, where $f'(u)$ denotes the first derivative of $f(u)$. Essentially, we approximate $\widehat{\psi}_{l-1}(\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z})$ in the partial likelihood function 3.3 by the first-order Taylor expansion at $\widehat{\boldsymbol{\beta}}_{l-1}$, which results in a concave function of $\boldsymbol{\beta}$. This optimization is solved through the Lagrange multiplier method. Given $\widehat{\boldsymbol{\beta}}_l$, we update $\boldsymbol{\alpha}$ and $\boldsymbol{\gamma}$ by maximizing $pl(\boldsymbol{\alpha}, \widehat{\boldsymbol{\beta}}_l, \boldsymbol{\gamma})$ under the constraint $-M_n \leqslant \gamma_1 \leqslant \cdots \leqslant \gamma_{K_n+1} \leqslant M_n$. Note that the objective function in this optimization is strictly concave and that the constraint sets are convex. Therefore, there exists a unique global maximum, and many software packages for convex optimization can be used. In particular, we adopt the R package "quadprog" (Berwin and Weingessel, 2013). We iterate through the above steps until convergence.

We need to determine K_n, M_n, and the knots for the B-splines. The asymptotic theory in the next section suggests $K_n = O(n^{1/13})$ and $M_n = O((\log n)^{\delta})$ for some $\delta \in (0, 1)$. In our experience, the estimates remain unchanged if M_n is set to be larger than 20; the performance is satisfactory when K_n ranges from 3 to 9; and K_n is determined by the AIC criterion (Akaike, 1970). For the boundary and interior knots, we use the extreme values and the $(K_n - 2)$ quantiles of $\{\widehat{\boldsymbol{\beta}}_{l-1}^{\mathrm{T}}\boldsymbol{Z}_i, i = 1, \cdots, n\}$, respectively, in each iteration in order to avoid sparse data when estimating the coefficients for B-splines.

Let $\widehat{\boldsymbol{\alpha}}$, $\widehat{\boldsymbol{\beta}}$, and $\widehat{\boldsymbol{\gamma}}$ be the final estimates of $\boldsymbol{\alpha}$, $\boldsymbol{\beta}$, and $\boldsymbol{\gamma}$, respectively. The link function $\psi(u)$ is estimated by $\widehat{\psi}(u) \equiv \widehat{\boldsymbol{\gamma}}^{\mathrm{T}}\boldsymbol{B}(u)$, and the cumulative baseline hazard function $\Lambda(\cdot)$ is estimated by the Breslow-type estimator (Breslow, 1972)

$$\sum_{i=1}^{n} \frac{I(\widetilde{T}_i \leqslant t)\Delta_i}{\sum_{j=1}^{n} Y_j(\widetilde{T}_i) \exp\{\widehat{\boldsymbol{\alpha}}^{\mathrm{T}}\boldsymbol{X}_j + A_j\widehat{\psi}(\widehat{\boldsymbol{\beta}}^{\mathrm{T}}\boldsymbol{Z}_j)\}}.$$

Finally, the optimal treatment rule is estimated as $I\{\widehat{\boldsymbol{\beta}}^{\mathrm{T}}\boldsymbol{Z} < \widehat{\psi}^{-1}(0)\}$. In addition, we can estimate the average treatment effect β_{av}, defined as the average log-hazard ratio of treatment A over all patients in the population $E\{\psi(\boldsymbol{\beta}^T\boldsymbol{Z})\}$, by $\widehat{\beta}_{\mathrm{av}} \equiv n^{-1}\sum_{i=1}^{n}\widehat{\psi}(\widehat{\boldsymbol{\beta}}^T\boldsymbol{Z}_i)$.

3.3 Asymptotic Theory

Let $\boldsymbol{\alpha}_0 \equiv (\alpha_{0,1}, \cdots, \alpha_{0,p})^{\mathrm{T}}$, $\boldsymbol{\beta}_0 \equiv (\beta_{0,1}, \cdots, \beta_{0,q})^{\mathrm{T}}$, and $\psi_0(\cdot)$ denote the true values of $\boldsymbol{\alpha}, \boldsymbol{\beta}$, and $\psi(\cdot)$, respectively. In addition, let τ denote the study duration and $\mathcal{Z}$ denote the union of the support of $\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z}$ for all $\|\boldsymbol{\beta}\| = 1$. We impose the following regularity conditions.

Condition 1. The true values $\boldsymbol{\alpha}_0$ and $\boldsymbol{\beta}_0$ lie in the interior of a known compact set; $\psi_0(u)$ is a strictly increasing function of u and three-times differentiable in $\mathcal{Z}$, and $\lambda_0(t)$ is continuously differentiable in $[0, \tau]$.

Condition 2. The conditional density of C given $\boldsymbol{X}$ is continuously differentiable on its support. The conditional distribution of $(A, \boldsymbol{X})$ given $\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z}$ has a continuously differentiable density with respect to a dominating measure.

Condition 3. With probability 1, $P(Y(\tau) = 1|\boldsymbol{X}) > c_0$ for some positive constant c_0. In addition, if $\boldsymbol{c}_1^{\mathrm{T}}\boldsymbol{X} = d_1$ with probability 1 for some constant vector $\boldsymbol{c}_1$ and constant d_1, then $\boldsymbol{c}_1 = \boldsymbol{0}$ and $d_1 = 0$.

Condition 4. If for some $\boldsymbol{\beta}^*$ with norm 1, $\mathrm{Var}(\boldsymbol{\beta}^{*\mathrm{T}}\boldsymbol{Z}|\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z}) = 0$ almost surely, then $\boldsymbol{\beta}^* = \pm\boldsymbol{\beta}_0$. In addition, $\partial/\partial\boldsymbol{\beta} E\{\psi_0(\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z})|\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z}\}\big|_{\boldsymbol{\beta}=\boldsymbol{\beta}_0}$ is non-degenerate.

Condition 5. The number of the knots satisfies that $K_n \to \infty$ and $n^{-1/2}K_n^7 \to \infty$. The upper bound M_n satisfies that $M_n \to \infty$, $n^{-1/2}K_n^4 M_n$ $\exp\{4M_n\} \to 0$ and $K_n^{-1/2}\exp\{M_n\} \to 0$. In addition, the adjacent distance of the interior knots is between $c^{-1}K_n^{-1}$ and cK_n^{-1} for some positive constant c.

Remark 1. The second part of Condition 1 and Condition 2 ensure smoothness for the functions ψ_0 and λ_0, the conditional distribution of C given $\boldsymbol{X}$, and the conditional distribution of $(A, \boldsymbol{X})$ given $\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z}$. Condition 3 ensures that a non-trivial proportion of subjects is censored at τ and that $(1, \boldsymbol{X}^{\mathrm{T}})^{\mathrm{T}}$ is linearly independent. Condition 4 is an identifiability condition used in single-index models. In particular, this condition holds if $\boldsymbol{Z}$ follows a multivariate normal distribution. In Condition 5, we may set K_n and M_n to $O(n^{1/13})$ and $O((\log n)^{\delta})$, respectively, for any constant $\delta \in (0, 1)$.

We state the consistency and asymptotic distribution of the estimators for the model parameters in the following three theorems, whose proofs are given in the Appendix. In addition, we establish the asymptotic properties for the estimator of the average treatment effect in §S.1 of the Supplementary Materials.

Theorem 3.3.1. *Under Conditions 1–5,* $||\widehat{\psi} - \psi_0||_{W^{1,\infty}(\mathcal{Z})} \to 0$, $||\widehat{\boldsymbol{\alpha}} - \boldsymbol{\alpha}_0|| \to 0$, *and* $||\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0|| \to 0$ *in probability, where for any differentiable function* f *with derivative* f', $||f||_{W^{1,\infty}(\mathcal{Z})}$ *is defined as* $||f||_{L^\infty(\mathcal{Z})} + ||f'||_{L^\infty(\mathcal{Z})}$. *Furthermore,* $||\widehat{\boldsymbol{\alpha}} - \boldsymbol{\alpha}_0||^2 + ||\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0||^2 + ||\widehat{\psi} - \psi_0||^2_{L_2(\mathcal{Z})} = o_p(n^{-1/2})$.

To describe the asymptotic distribution, we assume $\beta_{0q} > 0$ without loss of generality. For a q-dimensional vector $\boldsymbol{x}$, let $\boldsymbol{x}_{-q} = (x_1, \cdots, x_{q-1})^{\mathrm{T}}$. We introduce a $(p+q-1)$-vector function, $\boldsymbol{R}(s)$, as the solution to the following functional equation

$$\begin{aligned}\boldsymbol{R}(s) = E\Big\{\int AdN(u)|\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z} = s\Big\}^{-1} \\ \Big\{E\Big[\int E\{A\boldsymbol{R}(\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z})|T = u, \Delta = 1\}AdN(u)|\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z} = s\Big] + \\ E\Big(\int [\boldsymbol{g} - E\{\boldsymbol{g}(\boldsymbol{X}, A)|T = u, \Delta = 1\}]\, AdN(u)du|\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z} = s\Big)\Big\}, \quad (3.4)\end{aligned}$$

where $N(t) = \Delta I(T \leqslant t)$, and

$$\boldsymbol{g}(\boldsymbol{X}, A) = \begin{pmatrix} \boldsymbol{X} \\ A\psi_0'(\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z})\boldsymbol{Z}_{-q} - \boldsymbol{\beta}_{0,-q}A\psi_0'(\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z})Z_q/\beta_{0,q} \end{pmatrix}. \quad (3.5)$$

There exists a unique solution, as proven in §S.2 of the Supplementary Materials. Define

$$\begin{aligned}\boldsymbol{H}(u, \boldsymbol{X}, A) = \boldsymbol{g}(\boldsymbol{X}, A) - E\{\boldsymbol{g}(\boldsymbol{X}, A)|T = u, \Delta = 1\} - A\boldsymbol{R}(\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z}) \\ + E\{A\boldsymbol{R}(\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z})|T = u, \Delta = 1\}.\end{aligned}$$

Theorem 3.3.2. *Under Conditions 1–5,* $n^{1/2}(\widehat{\boldsymbol{\alpha}}^{\mathrm{T}} - \boldsymbol{\alpha}_0^{\mathrm{T}}, \widehat{\boldsymbol{\beta}}_{-q}^{\mathrm{T}} - \boldsymbol{\beta}_{0,-q}^{\mathrm{T}})^{\mathrm{T}}$ *converges in dis-*

tribution to a zero-mean normal random vector with covariance matrix $\boldsymbol{\Sigma}^{-1}$*, where* $\boldsymbol{\Sigma} = \int E\big[\boldsymbol{H}^{\otimes 2}(u, \boldsymbol{X}, A)\exp\{\boldsymbol{\alpha}_0^{\mathrm{T}}\boldsymbol{X} + A\psi_0(\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z})\}Y(u)\big]\lambda_0(u)du$*, and* $\boldsymbol{v}^{\otimes 2} = \boldsymbol{v}^{\mathrm{T}}\boldsymbol{v}$ *for any vector* $\boldsymbol{v}$*. In addition,* $\boldsymbol{\Sigma}^{-1}$ *achieves the semiparametric efficiency bound.*

To estimate $\boldsymbol{\Sigma}$, we need to estimate $\boldsymbol{R}$. To this end, let $\mathbb{P}$ and $\mathbb{P}_n$ denote the probability measure and the empirical measure, respectively. That is, for random variable U, $\mathbb{P}(U)$ is the expectation of U, and $\mathbb{P}_n(U)$ is the sample average over $\{U_i\}_{i=1}^n$. We partition the data into m_n groups based on the quantiles of $\{\widehat{\boldsymbol{\beta}}^{\mathrm{T}}\boldsymbol{Z}_i\}_{i=1}^n$, and we use $\widehat{E}_n(U|\widehat{\boldsymbol{\beta}}^{\mathrm{T}}\boldsymbol{Z} = s)$ to denote the average of $\{U_i\}_{i=1}^n$ among subjects who are in the same partition as s. In addition, define

$$\widetilde{E}_n[U] = \mathbb{P}_n[Y(u)\exp\{\widehat{\boldsymbol{\alpha}}^{\mathrm{T}}\boldsymbol{X} + A\widehat{\psi}(\widehat{\boldsymbol{\beta}}^{\mathrm{T}}\boldsymbol{Z})\}]^{-1}\mathbb{P}_n[UY(u)\exp\{\widehat{\boldsymbol{\alpha}}^{\mathrm{T}}\boldsymbol{X} + A\widehat{\psi}(\widehat{\boldsymbol{\beta}}^{\mathrm{T}}\boldsymbol{Z})\}].$$

In light of equation 3.4, we estimate $\boldsymbol{R}(s)$ by $\widehat{\boldsymbol{R}}(s)$, which solves the following equation

$$\begin{aligned}\widehat{\boldsymbol{R}}(s) =& \widehat{E}_n\Big\{\int AdN(u)|\widehat{\boldsymbol{\beta}}^{\mathrm{T}}\boldsymbol{Z} = s\Big\}^{-1}\Big\{\widehat{E}_n\Big[\int \widetilde{E}_n\{A\widehat{\boldsymbol{R}}(\widehat{\boldsymbol{\beta}}^{\mathrm{T}}\boldsymbol{Z})\}AdN(u)|\widehat{\boldsymbol{\beta}}^{\mathrm{T}}\boldsymbol{Z} = s\Big] \\ &+ \widehat{E}_n\Big(\int \Big[\boldsymbol{g} - \widetilde{E}_n\{\boldsymbol{g}(\boldsymbol{X}, A)\}\Big]AdN(u)du|\widehat{\boldsymbol{\beta}}^{\mathrm{T}}\boldsymbol{Z} = s\Big)\Big\}.\end{aligned}$$

Essentially, $\widehat{\boldsymbol{R}}(\cdot)$ solves a linear equation system. We then estimate $\boldsymbol{H}$ by

$$\widehat{\boldsymbol{H}}(u, \boldsymbol{X}, A) \equiv \boldsymbol{g}(\boldsymbol{X}, A) - \widetilde{E}_n\{\boldsymbol{g}(\boldsymbol{X}, A)\} - A\widehat{\boldsymbol{R}}(\widehat{\boldsymbol{\beta}}^{\mathrm{T}}\boldsymbol{Z}) + \widetilde{E}_n\{A\widehat{\boldsymbol{R}}(\widehat{\boldsymbol{\beta}}^{\mathrm{T}}\boldsymbol{Z})\}$$

and $\boldsymbol{\Sigma}$ by

$$\widehat{\boldsymbol{\Sigma}} \equiv \int \mathbb{P}_n\big[\widehat{\boldsymbol{H}}^{\otimes 2}(u, A, \boldsymbol{X})\exp\{\widehat{\boldsymbol{\alpha}}^{\mathrm{T}}\boldsymbol{X} + A\widehat{\psi}(\widehat{\boldsymbol{\beta}}^{\mathrm{T}}\boldsymbol{Z})\}Y(u)\big]d\widehat{\Lambda}(u).$$

Theorem 3.3.3. *Under Conditions 1–5, with* $m_n = cn^{1/2}$ *for a positive constant* c*,* $\widehat{\boldsymbol{\Sigma}} \to \boldsymbol{\Sigma}$ *in probability.*

3.4 Simulation Studies

3.4.1 Simulation results and comparison with other methods

We conducted extensive simulation studies to assess the performance of the proposed methods. We considered sample sizes $n = 500$, 1000, and 2000. We let the baseline hazard function $\lambda(t)$ follow a Weibull distribution with shape parameter 2.5 and scale parameter 2. We considered three link functions: (a) (exponential) $\psi_0(u) = e^u - 0.5$; (b) (linear) $\psi_0(u) = u + 0.4$; and (c) (sine)

$$\psi_0(u) = \begin{cases} -3 \text{ if } u \leqslant -1/2, \\ 2\sin(\pi u) - 1 \text{ if } -1/2 < u \leqslant 1/2, \\ 1 \text{ if } u > 1/2. \end{cases}$$

When the link function is linear, the proposed model reduces to the standard Cox model.

We let the censoring time follow Unif$[0, \tau]$, where τ was chosen to yield the censoring rate of 50%. We generated four independent covariates from Unif$[-1, 1]$ for the main effects and let the treatment effect on the survival time depend on the first two covariates. We set the main-effect parameters $(\alpha_1, \alpha_2, \alpha_3, \alpha_4)^{\mathrm{T}}$ to $(-0.4, -0.2, 0.2, 0.4)^{\mathrm{T}}$ and the interaction-effect parameters $(\beta_1, \beta_2)^{\mathrm{T}}$ to $(-0.6, 0.8)^{\mathrm{T}}$. We let the treatment assignment A be independent of $\boldsymbol{X}$ and follow Bern(0.5). We simulated 10,000 replicates for each setting.

We considered a series of models with the number of knots K_n ranging from 3 to 9. The majority of the 10,000 replicates selected a final model with $K_n = 3, 3$, and 6 by the AIC criterion for the exponential, linear, and sine link functions, respectively. The estimates remained unchanged with an M_n larger than 20. For $n = 500$, 99.4% of the replicates converged with 500 iterations; for $n = 2000$, the convergence rate was higher than 99.95%. It took approximately 3 seconds and 2 minutes to analyze one simulated dataset for $n = 500$ and $n = 2000$, respectively.

Table 1 summarizes the results for the estimation of $(\beta_1, \beta_2)^{\mathrm{T}}$ under the proposed model.

Table 3.1: Simulation results for the estimation of $\boldsymbol{\beta}$

		Proposed model				Cox model			
$\boldsymbol{n}$		**Bias**	**SE**	**SEE**	**CP**	**Bias**	**SE**	**SEE**	**CP**
Exponential link function									
500	β_1	0.009	0.147	0.151	0.926	0.011	0.154	0.155	0.923
	β_2	-0.015	0.115	0.114	0.918	-0.015	0.117	0.118	0.907
1000	β_1	0.004	0.100	0.105	0.943	0.005	0.105	0.106	0.934
	β_2	-0.007	0.076	0.079	0.940	-0.007	0.079	0.080	0.929
2000	β_1	0.001	0.070	0.073	0.951	0.000	0.074	0.074	0.939
	β_2	-0.004	0.053	0.055	0.948	-0.005	0.056	0.056	0.937
Linear link function									
500	β_1	0.026	0.207	0.183	0.886	0.018	0.184	0.189	0.911
	β_2	-0.024	0.159	0.136	0.864	-0.020	0.139	0.147	0.892
1000	β_1	0.012	0.138	0.129	0.919	0.009	0.126	0.128	0.931
	β_2	-0.010	0.103	0.096	0.906	-0.009	0.095	0.097	0.921
2000	β_1	0.004	0.093	0.090	0.934	0.003	0.088	0.089	0.939
	β_2	-0.005	0.070	0.067	0.929	-0.005	0.067	0.067	0.935
Sine link function									
500	β_1	0.002	0.042	0.043	0.957	0.071	0.083	0.084	0.883
	β_2	0.000	0.031	0.032	0.952	0.043	0.052	0.053	0.803
1000	β_1	0.001	0.028	0.029	0.960	0.069	0.057	0.059	0.812
	β_2	0.000	0.021	0.022	0.959	0.045	0.036	0.037	0.727
2000	β_1	0.001	0.019	0.020	0.958	0.069	0.041	0.042	0.628
	β_2	0.001	0.015	0.015	0.958	0.046	0.025	0.026	0.549

Bias and SE are the bias and standard error of the parameter estimator, respectively; SEE is the mean of the standard error estimator; CP is the coverage probability of the 95% confidence interval. The SEE and CP in the Cox model are based on the robust variance estimator.

For comparison, we also show the estimation results for a normalized estimator of the treatment-covariate interaction under the standard Cox model. The biases of the parameter estimators under our model are small and decrease as n increases. The variance estimators and the corresponding confidence intervals for $\boldsymbol{\beta}$ become more accurate as n increases. For simulation settings with a small sample size, the bootstrap method may provide a more accurate variance estimator and confidence interval; example simulations are included in section A.2 of the appendix. In the case of the sine link function, whose first-order derivative has large fluctuations over the values of the linear predictor $\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z}$, the estimation of the

interaction effects under the standard Cox model is highly biased.

Figure 1 shows that the link function is estimated accurately by the proposed methods. Table 2 summarizes the results for estimating the average treatment effect. The estimator of the average treatment effect is nearly unbiased. The variance estimator is accurate and the corresponding confidence intervals have correct coverages, especially for large n.

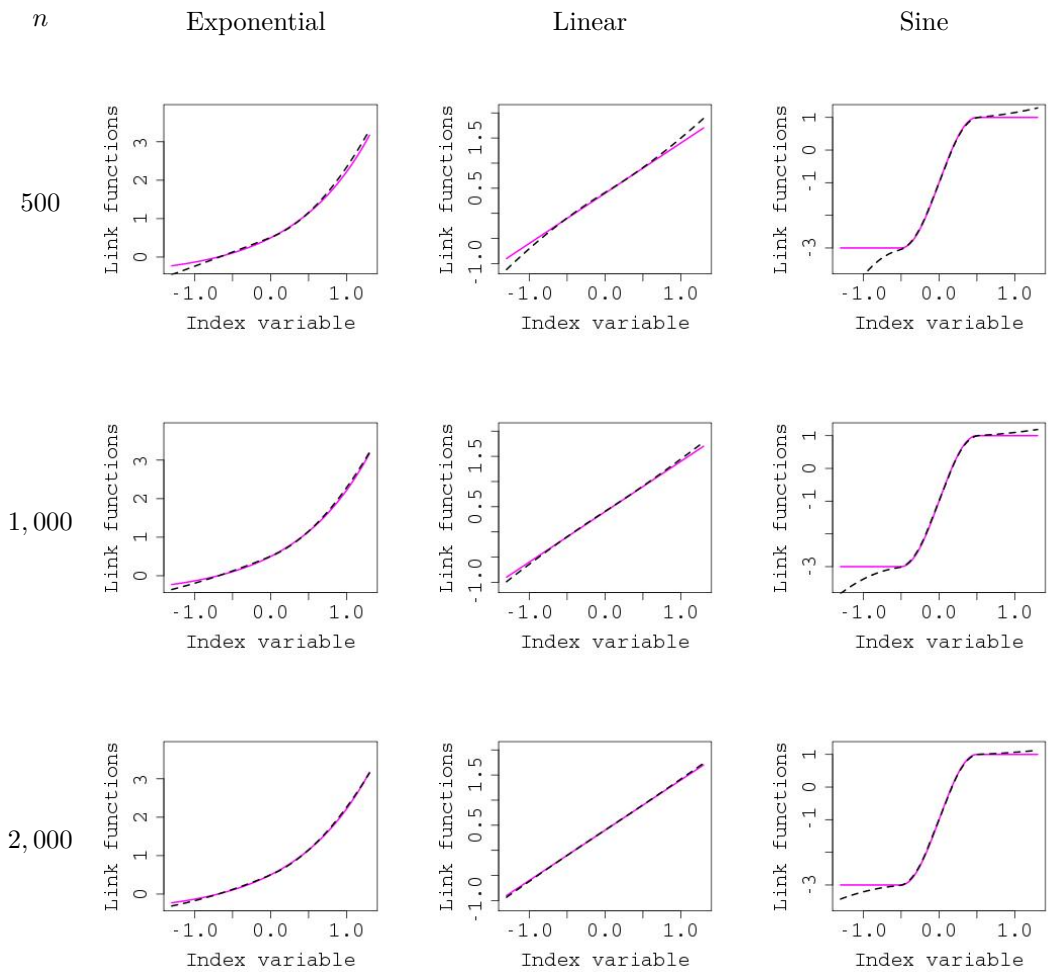

Figure 3.1: Estimation of the link function ψ. The solid and dotted curves pertain to the true value and the mean estimate, respectively.

In order to assess the performance of the treatment rules, we generated an independent "test" dataset and used the estimated coefficients under the proposed model, the Cox model, and the method of Jiang et al. (2017) to derive the optimal treatment rules for each subject.

In particular, we used $t = 2$ and the SAIPSWKME method that is described in the paper by Jiang et al. (2017), referred to as "Jiang's method" in the rest of the paper. We calculated the correct treatment assignment rate and the average survival probability of patients who follow the treatment rules under each model. For each patient in the test dataset, we compared the treatment assignment recommended by each model with the optimal treatment rule. Table 3 shows the average treatment assignment rates using the sine link function and $n = 1,000$. The proposed model has an incorrect treatment assignment rate of 1.6%, whereas the Cox model and Jiang's method have higher incorrect assignment rates of 13.2% and 4.9%, respectively. The Cox model also tends to recommend treatment 1 to patients whose true optimal treatment is 0.

Table 3.2: Simulation results for the average treatment effect

Link	$\boldsymbol{n}$	**Bias**	**SE**	**SEE**	**CP**
Exponential	500	0.015	0.141	0.145	0.954
	1000	0.010	0.098	0.098	0.952
	2000	0.006	0.068	0.069	0.952
Linear	500	0.011	0.139	0.143	0.954
	1000	0.008	0.096	0.097	0.952
	2000	0.005	0.067	0.067	0.954
Sine	500	-0.085	0.264	0.265	0.921
	1000	-0.030	0.146	0.146	0.953
	2000	-0.014	0.096	0.098	0.960

See the Note to Table 3.1.

Table 3.3: Simulation results on the treatment assignments

Optimal treatment	**Proposed model trt. assign.**		**Cox model trt. assign.**		**Jiang's method trt. assign.**	
	0	1	0	1	0	1
0	38.2%	0.8%	25.9%	13.1%	35.8%	3.2%
1	0.8%	60.2%	0.1%	60.9%	1.7%	59.3%

Parameters in all models were estimated in simulation scenario (c): sine link function. The correct treatment assignment rate was estimated from a test dataset of 1,000 subjects.

For each patient in the test dataset, we calculated the true survival functions corresponding

to the treatment rules determined by the proposed model and by the Cox model. Patients were divided into three subgroups according to their "risk score" $\mathcal{SC}(\boldsymbol{Z}) \equiv 0.8Z_2 - 0.6Z_1$. Figure 2 shows the average survival probability over time for each subgroup of patients in the case of the sine link function with $n = 1,000$. For subgroups (a) and (c) in Figure 2, the survival curves are similar between the treatment rule determined by the proposed model and that determined by the Cox model.

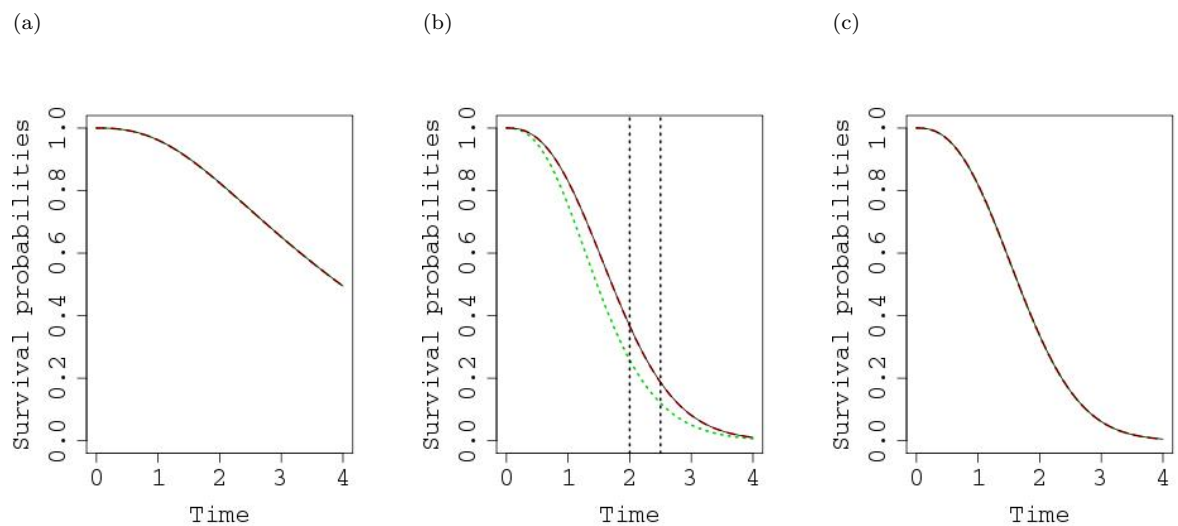

Figure 3.2: Estimation of the survival functions for patients with: (a) $\mathcal{SC}(\boldsymbol{Z}) < 0.2$; (b) $0.2 \leqslant \mathcal{SC}(\boldsymbol{Z}) \leqslant 0.5$; and (c) $\mathcal{SC}(\boldsymbol{Z}) > 0.5$. The dashed, dotted, and solid curves pertain to the survival functions under the proposed model, the Cox model, and the optimal treatment assignment, respectively.

For the 183 subjects with risk scores between 0.2 and 0.5 in subgroup (b), Figure 2 shows that the average survival probability over time is much higher if patients follow the treatments assigned by the proposed model as opposed to those assigned by the Cox model. Specifically, at time $t = 2$, patients with $\mathcal{SC}(\boldsymbol{Z}) \in [0.2, 0.5]$ who receive the treatments recommended by the proposed model have an average survival probability of 0.37, which is the same as that of the optimal treatment rule. By contrast, if they receive the treatments recommended by the Cox model, the average survival probability is only 0.26. In other words, the adoption

of the treatment rule formulated by the proposed model improves average patient survival probability by 42% among patients with $\mathcal{SC}(\boldsymbol{Z}) \in [0.2, 0.5]$, and the proposed model performs as well as the true optimal treatment assignment.

3.4.2 Treatment assignment under model misspecification

We compared the estimated treatment regimen under the proposed single-index model with the Cox model and Jiang's method under various model misspecifications. We first considered I) the proportional odds model, where assumptions in Jiang's method still hold, but the proposed model and the Cox model are misspecified. We simulated the data with true values specified as in the previous subsection with an exponential link and sample size of 1000. We further considered the following two cases: II) when the baseline covariate effect is non-linear and assumptions in Jiang's method still hold, but the proposed model and the Cox model are misspecified; and III) when the interaction effect is captured by the sum of two indices and all three models are misspecified. Specifically, for scenario II), we considered the hazard-ratio relationship

$$\lambda(t; \boldsymbol{X}, A) = \lambda_0(t) \exp\left\{\psi_1(\boldsymbol{\alpha}^{\mathrm{T}}\boldsymbol{X}) + A\psi_2(\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z})\right\}.$$

Let $\psi_1(u) = u^3 + 0.6$, and let ψ_2 take the sine link form, as in subsection 4.1. For scenario IV), we considered the hazard-ratio relationship

$$\lambda(t; \boldsymbol{X}, A) = \lambda_0(t) \exp\left\{\boldsymbol{\alpha}^{\mathrm{T}}\boldsymbol{X} + cA\psi_1(\boldsymbol{\beta}_1^{\mathrm{T}}\boldsymbol{Z}) + A\psi_2(\boldsymbol{\beta}_2^{\mathrm{T}}\boldsymbol{Z})\right\},$$

where $\psi_1(u) = u^3 + 0.6$, $\psi_2(u) = e^u - 0.5$, and $c \in \{0.2, 0.6\}$. We set the interaction-effect parameters to $\beta_2 = (-0.91, -0.41)$ and $\beta_1 = (-0.6, 0.8)$. Other aspects of the simulations in II) and III) were the same as in subsection 4.1. We simulated 1000 replicates for all settings.

Table 4 summarizes the treatment assignments under the three models versus the true optimal treatment under different model misspecifications. In setting I), the Cox model shows the smallest misclassification rate of 8.78%, and the proposed model is the second best, with

a misclassification rate of 9.35%. The proposed model gives the most accurate treatment rule, with a smaller misclassification rate in settings II) and III).

Table 3.4: Simulation results on the treatment assignment with different hazards relationships

Optimal trt.	**Proposed model trt. assgn.**		**Cox model trt. assgn.**		**Jiang's method trt. assgn.**	
	0	1	0	1	0	1
I) Proportional Odds Models						
0	81.91%	5.13%	79.45%	7.60%	76.77%	10.27%
1	4.22%	8.73%	1.18%	11.78%	4.74%	8.21%
II) Cubic Main Effect						
0	38.72%	0.97%	11.22%	28.46%	14.13%	25.55%
1	0.75%	59.57%	16.94%	43.38%	21.29%	39.03%
III) Summation of two single index terms for interaction with $c = 0.2$						
0	86.67%	3.32%	74.11%	15.89%	77.59%	12.40%
1	2.57%	7.44%	8.23%	1.78%	8.61%	1.39%
III) Summation of two single index terms for interaction with $c = 0.6$						
0	89.23%	2.25%	77.65%	13.83%	81.94%	9.54%
1	1.88%	6.65%	7.22%	1.30%	7.62%	0.90%

In the table, "trt. assgn." stands for treatment assignment rate, which was estimated from a test dataset of 1000 patients.

3.5 Application to an AIDS clinical trial

The ACTG 175 study is a randomized, double-blind, phase II/III clinical trial (Hammer et al., 1996). A total of 2,139 subjects infected with human immunodeficiency virus type 1 (HIV-1) were randomized to one of four arms: zidovudine (ZDV) monotherapy; didanosine (ddI) monotherapy; ZDV + ddI; or ZDV + zalcitabine (ddC). The primary endpoint was $\geqslant$ 50% decline in CD4 cell count, development of AIDS, or death. Approximately 75.6% of the subjects were censored.

We aimed to identify baseline covariates that influence the treatment effects on the survival time and derive the optimal treatment regime for each patient in the study. Following Jiang et al. (2017), we only considered the 1,054 subjects who received ZDV + ddI (A = 1) or ZDV + ddC (A = 0), and we included the following 12 covariates for the main effects: age;

Table 3.5: Estimation of the parameters under the proposed, Cox and Jiang's models

	Proposed model			**Cox model**		
Parameters	**Est.**	**Std. Err.**	***P*-value**	**Est.**	**Std. Err.**	***P*-value**
Main Effects						
Weight	0.16	0.10	0.11	0.14	0.10	0.17
Age	0.07	0.10	0.44	0.06	0.10	0.54
CD4 count	−0.56	0.12	< 0.01	−0.59	0.12	< 0.01
CD8 count	0.16	0.07	0.03	0.16	0.07	0.02
Karnofsky score	−0.36	0.09	< 0.01	−0.32	0.09	< 0.01
Antiretroviral history	0.46	0.16	< 0.01	0.46	0.16	< 0.01
Gender	−0.18	0.32	0.58	−0.18	0.31	0.56
Race	−0.07	0.18	0.71	−0.07	0.18	0.71
Homosexual activity	0.26	0.27	0.35	0.26	0.27	0.33
Intravenous drug use	−0.58	0.25	0.02	−0.57	0.25	0.02
Hemophilia status	0.28	0.35	0.41	0.29	0.34	0.39
Symptomatic status	0.52	0.17	< 0.01	0.52	0.16	< 0.01
Average treatment effect	0.01	0.16	0.96	−0.12	0.14	0.40
Interactions						
Age	−0.84	0.15	< 0.01	−0.81	0.23	< 0.01
Karnofsky score	0.42	0.25	0.09	0.27	0.35	0.44
Weight	0.08	0.27	0.77	0.14	0.36	0.69
CD4 count	0.33	0.29	0.26	0.50	0.37	0.18

The estimate of the interaction effect in the Cox model is scaled to have the same Euclidean norm as that in the proposed model. "Est." and "Std. Err." denote the estimate and the standard error estimate for a parameter, respectively.

weight; Karnofsky score; baseline CD4 cell count; baseline CD8 cell count; hemophilia status; homosexual activity; history of intravenous drug use; race; gender; antiretroviral history; and symptomatic status. The results in Jiang et al. (2017) and Geng et al. (2015) suggested that Karnofsky score, weight, CD4 cell count, and age may interact with treatment effects. These four covariates were included in our model as interaction terms. We characterized the treatment-covariate interaction effects on time to death and derived the treatment rules by fitting model 2.1. All covariates were standardized.

We fit model 2.1 with the number of knots K_n ranging from 3 to 7 and set M_n to 20. By the AIC criterion, we selected the model with $K_n = 3$ and the estimated link function in

Figure 3. We checked the goodness-of-fit of the model by examining the martingale residuals, which is discussed in details in Section S.5 in the supplementary materials. Table 5 displays the estimation results for the average treatment effect and the treatment-covariate interaction effects under the proposed model and the Cox model. Both models indicate that age is significantly related to the treatment effect on patient survival, with older patients more likely to benefit from treatment ZDV + ddI than treatment ZDV + ddC. The Karnofsky score plays a more important role in the optimal treatment rule estimation under the proposed model than under the Cox model. The average treatment effect is not significant under either the proposed model or the Cox model.

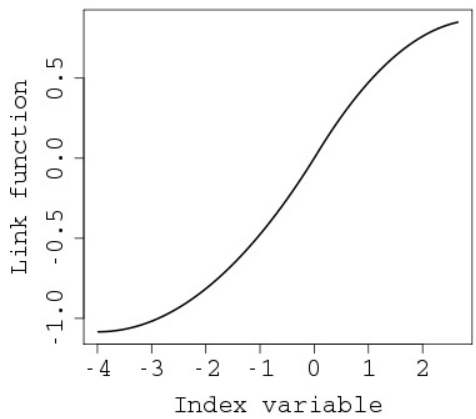

Figure 3.3: Estimation of the link function ψ.

We compared the treatment rule estimated by the proposed model with the Cox model and Jiang's method at $t = 400$ using SIAPSWKME. The time $t = 400$ was used in the analysis of the same dataset in Jiang et al. (2017). The treatment rule estimated by the proposed model, the Cox model and Jiang's method is, respectively, $\{-0.42 * \text{Karnofsky score} - 0.33 * \text{CD4 cell count} + 0.84 * \text{age} - 0.08 * \text{weight} < 0\}$, $\{-0.27 * \text{Karnofsky score} - 0.50 * \text{CD4 cell count} + 0.81 * \text{age} - 0.14 * \text{weight} < 0\}$, and $\{-0.24 * \text{Karnofsky score} + 0.40 * \text{CD4 cell count} + 0.58 * \text{age} - 0.42 * \text{weight} + 0.52 < 0\}$. Table 6 shows that, out of the 1,046 patients, 74 would receive different treatments under the proposed model

versus the Cox model, 375 would receive different treatments under the proposed model versus Jiang's model. Figure 4 shows that the survival curves are comparable in the two treatment arms, with the survival probabilities improved if patients follow the treatment regimen recommended by the proposed model.

Table 3.6: Treatment assignments under the proposed and Cox models

	Cox model		**Jiang's model**		
Proposed model	**ZDV+ ddI**	**ZDV+ ddC**	**ZDV+ ddI**	**ZDV+ ddC**	**Total**
ZDV + ddI	458	26	437	47	484
ZDV + ddC	48	514	328	234	562
Total	506	540	765	281	1046

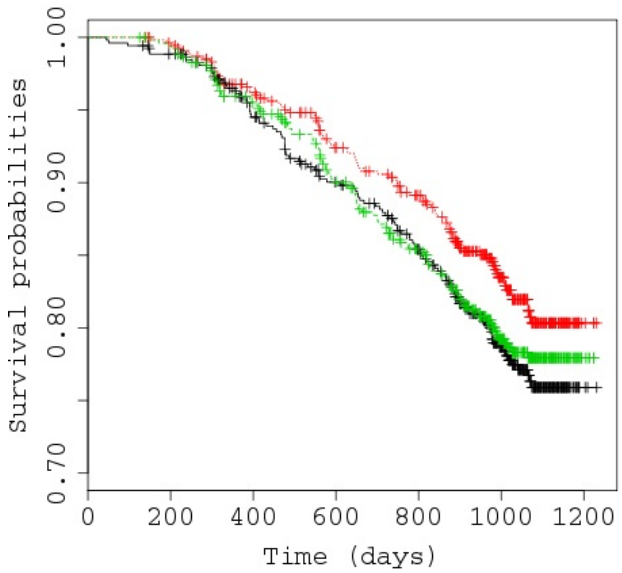

Figure 3.4: Estimation of the survival function for the ACTG 175 study. The black, green, and red curves pertain, respectively, to the ZDV + ddC arm, ZDV + ddI arm, and those patients who would have received treatments assigned by the proposed model.

To assess how well each treatment regimen works, we estimated the survival functions for patients based on the assigned treatments and the opposite ones. These two groups

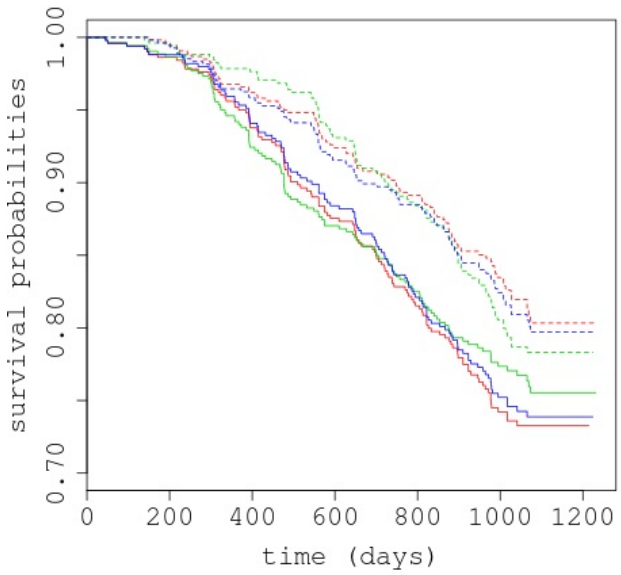

Figure 3.5: Estimation of the survival functions for patients in the ACTG 175 study assigned to groups classified as “agreed" (dashed line) or “disagreed" (solid line) based on treatment assignments. The red, blue, and greens curves pertain, respectively, to the proposed model, the Cox model, and Jiang’s model.

are referred to as the “agreed" case and the “disagreed" case, respectively. We performed a log-rank test to compare the survival functions between these two hypothetical groups of patients. Survival curves for the “agreed" group and the “disagreed" group are well separated under the proposed model (p-value = 0.002) but less so under the Cox model (p-value = 0.01) or Jiang’s model (p-value = 0.13); see Figure 5. In particular, the difference of estimated survival probabilities between the “agreed" and “disagreed" groups on day 600 is 0.031 under the Cox model and 0.049 under the proposed model (i.e., 60% higher for the proposed model). Survival probabilities difference is larger around $t = 400$ under Jiang’s method, as their method aims to maximize the survival probability difference at that time point. At later time points such as $t = 1000$, the curves are better separated under the proposed model than Jiang’s method. This finding suggests that the treatment rule formulated by the proposed model is more beneficial to patients, as measured by the survival function.

3.6 Discussion

In this chapter, we proposed a framework to estimate personalized treatment rules by modeling treatment-covariate interactions through a single index function. We developed an efficient and stable numerical algorithm based on approximate functions to overcome the non-convexity problem in the optimization step. We derived the asymptotic distributions for the estimators of the treatment-covariate interaction effects and the average treatment effect. The proofs are challenging due to the presence of two different non-parametric components in the model. The limiting covariance matrix involves a function that is the solution to an equation with no analytic form. By approximating this function with a step function and applying the kernel estimation technique, we transformed the original equation into a finite-dimensional linear system that can be solved via a one-step procedure. People may be interested in quantities that summarize the patients performance, such as the survival function under the treatment rule. These quantities could be estimated by the plug-in estimator, and their asymptotic properties follow from the delta method and the limiting distribution of $\widehat{\boldsymbol{\alpha}}$, $\widehat{\boldsymbol{\beta}}$ and $\widehat{\psi}$.

Several new research directions extend naturally from the work presented in this chapter. First, we will consider other types of models, such as linear transformation models (Zeng and Lin, 2007), in order to allow different types of dependency of the survival time on covariates and treatment. Second, we will develop variable-selection procedures to handle higher-dimensions of interaction effects. Finally, we will develop methods that allow for more than two treatment options and multiple decision points.

CHAPTER 4: TRANSFORMATION SINGLE-INDEX MODELS FOR OPTIMAL TREATMENT REGIMENS

4.1 Introduction

A flexible semiparametric single-index model was proposed in Chapter 3 to assess how the effect of treatment on the survival distribution depends on each individual covariate. The model was flexible enough to allow complex treatment-covariate interactions while providing an interpretable linear treatment rule. In scientific studies, for example when modeling prognostic factors with medium or long follow-up times, the proportional hazards assumption is often violated (Valsecchi et al., 1996). The class of linear transformation models (Dabrowska and Doksum, 1988), which includes both the proportional hazards and the proportional odds models as special cases, could serve as an useful alternative. In order to accommodate general hazard relationship, it is important to extend the single-index proportional hazards model to a more flexible framework of transformations. The nonconcave likelihood and the infinite-dimensional nuisance parameters in the single-index transformation model make the computation difficult. The partial likelihood approach used in the derivation of asymptotic properties for the single-index proportional hazards model no longer applies here, as the cumulative hazard function cannot be profiled out with the general transformation.

In this chapter, we propose a semiparametric single-index model in the transformation model framework to assess the individualized treatment effect. The proposed model is an extension of Chapter 3 by allowing general hazard relationships. The treatment rule derived from the proposed model remains linear, which is computationally simple and clinically interpretable for personalized treatment decisions. We show that by casting the optimization of likelihood with transformation function into the missing value framework, the EM algorithm could be used to simplify the estimation. We develop the associate asymptotic theory and

show that the resulting estimators are semiparametrically efficient. The performance of the proposed methods is illustrated in simulation studies and a real data application.

4.2 Methods

4.2.1 Data and Model

Let T, A and $\boldsymbol{X}$ denote the survival time, the binary treatment with values 0 and 1, and bounded baseline covariates, respectively. To characterize the heterogeneity of treatment effects under the semiparametric transformation model, we assume that the cumulative hazard function for T conditional on A and $\boldsymbol{X}$ takes the form

$$\Lambda(t; A, \boldsymbol{X}) = G\left[\int_0^t \exp\{\boldsymbol{\alpha}^{\mathrm{T}}\boldsymbol{X} + A\psi(\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z})\}d\Lambda(s)\right], \tag{4.1}$$

where $G(\cdot)$ is a specific transformation function that is strictly increasing, $\Lambda(\cdot)$ is an unknown increasing function (Zeng & Lin, 2006), $\boldsymbol{Z}$ is a subset of $\boldsymbol{X}$, $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ are the regression parameters, and $\psi(\cdot)$ is an unknown and strictly increasing link function. We assume that there are non-zero components in $\boldsymbol{\beta}$ and $||\boldsymbol{\beta}|| = 1$ for identifiability, where $||\boldsymbol{x}||$ is the Euclidean norm of vector $\boldsymbol{x}$.

Clearly, when $\psi(x) = x$, model (4.1) reduces to the usual transformation model with the interaction between the treatment and the covariates. The choice of $G(x) = x$ yields the proportional hazards models with a single index term for treatment covariate interaction. An important class of transformation models is the class of logarithmic transformations $G(x) = r^{-1}\log(1 + rx)(r \geqslant 0)$, with $r = 0$ and $r = 1$ corresponding to the proportional hazards and proportional odds models, respectively. Since all the covariates are time-independent, model (4.1) can be rewritten as a linear transformation model $\log\Lambda(T) = -\boldsymbol{\alpha}^{\mathrm{T}}\boldsymbol{X} - A\psi(\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z}) + \epsilon$, where ϵ is an error term with distribution function $1 - \exp[-G\{\exp(x)\}]$ (Chen et al., 2002).

4.2.2 Maximum Likelihood Estimation

Suppose that data from a randomized trial consist of n right-censored observations. We denote them as $(\widetilde{T}_i, \Delta_i, A_i, \boldsymbol{X}_i)$ $(i = 1, \ldots, n)$, where $\widetilde{T}_i = \min(T_i, C_i)$ and $\Delta_i = I(T_i \leqslant C_i)$,

and C_i is the potential censoring time for subject i. We further assume that C_i is independent of T_i given $(A_i, \boldsymbol{X}_i)$. The observed log-likelihood function concerning the model parameters is then given by

$$\sum_{i=1}^{n} \left(\Delta_i \left\{ \boldsymbol{\alpha}^{\mathrm{T}} \boldsymbol{X}_i + A_i \psi(\boldsymbol{\beta}^{\mathrm{T}} \boldsymbol{Z}_i) \right\} + \Delta_i \log \lambda(\widetilde{T}_i) + \Delta_i \log G' \left[\exp\{\boldsymbol{\alpha}^{\mathrm{T}} \boldsymbol{X}_i + A_i \psi(\boldsymbol{\beta}^{\mathrm{T}} \boldsymbol{Z}_i)\} \Lambda(\widetilde{T}_i) \right] \right.$$
$$\left. - G \left[\exp\{\boldsymbol{\alpha}^{\mathrm{T}} \boldsymbol{X}_i + A_i \psi(\boldsymbol{\beta}^{\mathrm{T}} \boldsymbol{Z}_i)\} \Lambda(\widetilde{T}_i) \right] \right), \quad (4.2)$$

where $\Lambda(t) = \int_0^t \lambda(s) ds$.

Expression (4.2) involves both $\boldsymbol{\alpha}, \boldsymbol{\beta}$ and the infinite dimensional parameters ψ and Λ, and it may not be concave in these parameters. Also, there is no partial likelihood function available due to the transformation G. To resolve these difficulties, we adopt a mixture of expectation-maximization (EM) algorithm, NPMLE, and sieve estimation approach for the estimation procedure. First, for all commonly used transformations, $\exp\{-G(x)\}$ is the Laplace transformation of some function $\psi(x)$ such that

$$\exp\{-G(x)\} = \int_0^\infty \exp(-xt)\phi(t)dt.$$

Also,

$$G'(x) \exp\{-G(x)\} = \int_0^\infty t \exp(-xt)\phi(t)dt.$$

As proposed in Zeng and Lin (2007), by introducing a new frailty ξ with density function ϕ, the objective function can be written as

$$\prod_{i=1}^{n} \int \left[\xi_i \exp\left\{ \boldsymbol{\alpha}^{\mathrm{T}} \boldsymbol{X}_i + A_i \boldsymbol{\gamma}^{\mathrm{T}} \boldsymbol{B}(\boldsymbol{\beta}^{\mathrm{T}} \boldsymbol{Z}_i) \right\} \Lambda\{\widetilde{T}_i\} \right]^{\Delta_i}$$
$$\exp\left[-\xi_i \exp\{\boldsymbol{\alpha}^{\mathrm{T}} \boldsymbol{X}_i + A_i \boldsymbol{\gamma}^{\mathrm{T}} \boldsymbol{B}(\boldsymbol{\beta}^{\mathrm{T}} \boldsymbol{Z}_i)\} \Lambda(\widetilde{T}_i) \right] \phi(\xi_i) d\xi_i. \quad (4.3)$$

This expression is the likelihood function under the proportional hazards frailty model with

conditional hazard function

$$\xi_i \lambda_0(t) \exp\left\{\boldsymbol{\alpha}^{\mathrm{T}} \boldsymbol{X}_i + A_i \boldsymbol{\gamma}^{\mathrm{T}} \boldsymbol{B}(\boldsymbol{\beta}^{\mathrm{T}} \boldsymbol{Z}_i)\right\}.$$

That is, we can use the EM algorithm in the estimation procedure by treating $\{\xi_i\}$ as missing data following a distribution with a (set of) parameter r.

If we could observe $\{\xi_i\}$, the full log-likelihood is, up to a term free of the parameter values,

$$L_{Full}(r, \boldsymbol{\alpha}, \boldsymbol{\beta}, \psi, \Lambda; \{\Delta_i, \widetilde{T}_i, \boldsymbol{X}_i, A_i, \xi_i\}) = L_1(r) + L_2(\boldsymbol{\alpha}, \boldsymbol{\beta}, \psi, \Lambda),$$

where $L_1(r)$ only involves the latent variable $\{\xi_i\}$ and parameter r, and

$$L_2(\boldsymbol{\alpha}, \boldsymbol{\beta}, \psi, \Lambda) = \sum_{i=1}^{n} \left[\Delta_i \left\{\boldsymbol{\alpha}^{\mathrm{T}} \boldsymbol{X}_i + A_i \psi(\boldsymbol{\beta}^{\mathrm{T}} \boldsymbol{Z}_i) + \log \lambda(\widetilde{T}_i)\right\} - \xi_i \Lambda(\widetilde{T}_i) \exp\{\boldsymbol{\alpha}^{\mathrm{T}} \boldsymbol{X}_i + A_i \psi(\boldsymbol{\beta}^{\mathrm{T}} \boldsymbol{Z}_i)\}\right].$$

In the E step, as shown in the appendix of Zeng and Lin (2007), the posterior expectation of the latent variables $\{\xi_i\}$, given the data and current estimates of all parameters $(\widetilde{\boldsymbol{\alpha}}, \widetilde{\boldsymbol{\beta}}, \widetilde{\psi}, \widetilde{\Lambda})$ can be expressed as

$$W_i \equiv E[\xi_i | \Delta_i, \widetilde{T}_i, \boldsymbol{X}_i, A_i, \widetilde{\boldsymbol{\alpha}}, \widetilde{\boldsymbol{\beta}}, \widetilde{\psi}, \widetilde{\Lambda}] = G'\left\{\exp\{\widetilde{\boldsymbol{\alpha}}^{\mathrm{T}} \boldsymbol{X}_i + A_i \widetilde{\psi}(\widetilde{\boldsymbol{\beta}}^{\mathrm{T}} \boldsymbol{Z}_i)\} \widetilde{\Lambda}(\widetilde{T}_i)\right\} - \Delta_i \frac{G''\left\{\exp\{\widetilde{\boldsymbol{\alpha}}^{\mathrm{T}} \boldsymbol{X}_i + A_i \widetilde{\psi}(\widetilde{\boldsymbol{\beta}}^{\mathrm{T}} \boldsymbol{Z}_i)\} \widetilde{\Lambda}(\widetilde{T}_i)\right\}}{G'\left\{\exp\{\widetilde{\boldsymbol{\alpha}}^{\mathrm{T}} \boldsymbol{X}_i + A_i \widetilde{\psi}(\widetilde{\boldsymbol{\beta}}^{\mathrm{T}} \boldsymbol{Z}_i)\} \widetilde{\Lambda}(\widetilde{T}_i)\right\}},$$

where $\Lambda\{t\}$ denotes the step size of Λ at t. The resulting expectation of L_2 given the data

and the current values of W_i is

$$\widetilde{L}_2(\boldsymbol{\alpha}, \boldsymbol{\beta}, \psi, \Lambda) = \sum_{i=1}^{n} \left[\Delta_i \left\{ \boldsymbol{\alpha}^{\mathrm{T}} \boldsymbol{X}_i + A_i \psi(\boldsymbol{\beta}^{\mathrm{T}} \boldsymbol{Z}_i) + \log \lambda(\widetilde{T}_i) \right\} - W_i \Lambda(\widetilde{T}_i) \exp\{\boldsymbol{\alpha}^{\mathrm{T}} \boldsymbol{X}_i + A_i \psi(\boldsymbol{\beta}^{\mathrm{T}} \boldsymbol{Z}_i)\} \right]. \quad (4.4)$$

The M step of the EM algorithm requires the maximization of 4.4 with respect to the unknown parameters $\boldsymbol{\alpha}$, $\boldsymbol{\beta}$, ψ, Λ. Since 4.4 is a weighted log likelihood of the Cox model, we can apply the computational algorithm in section 3.2.3. Specifically, we let the estimator for Λ be a step function with jumps only at the observed $\widetilde{T}_i$ with $\Delta_i = 1$. By maximizing 4.4 with respect to the jump sizes, Λ is profiled out, which yields

$$L_3(\boldsymbol{\alpha}, \boldsymbol{\beta}, \psi) = \sum_{i=1}^{n} \Delta_i \left(\boldsymbol{\alpha}^{\mathrm{T}} \boldsymbol{X}_i + A_i \psi(\boldsymbol{\beta}^{\mathrm{T}} \boldsymbol{Z}_i) - \log \left[\sum_{j=1}^{n} Y_j(\widetilde{T}_i) W_j \exp\{\boldsymbol{\alpha}^{\mathrm{T}} \boldsymbol{X}_j + A_j \psi(\boldsymbol{\beta}^{\mathrm{T}} \boldsymbol{Z}_j)\} \right] \right). \quad (4.5)$$

We use B-splines (Schumaker, 1981) to approximate the link function ψ. Specifically, we let $\boldsymbol{B}(u) = (B_1(u), \cdots, B_{K_n+1}(u))^{\mathrm{T}}$ denote quadratic B-spline bases corresponding to K_n distinct knots in an interval containing the support of $\boldsymbol{\beta}^{\mathrm{T}} \boldsymbol{Z}$ for any $\|\boldsymbol{\beta}\| = 1$. We then approximate $\psi(u)$ by $\widetilde{\psi}(u) \equiv \boldsymbol{\gamma}^{\mathrm{T}} \boldsymbol{B}(u)$, where $\boldsymbol{\gamma} = (\gamma_1, \cdots, \gamma_{K_n+1})^{\mathrm{T}}$ are unknown coefficients. To further ensure that $\widetilde{\psi}(u)$ is increasing but not too large, we require $-M_n \leqslant \gamma_1 \leqslant \cdots \leqslant \gamma_{K_n+1} \leqslant M_n$ for some constant M_n. The increasing sequence of γ's guarantees the increasing property of $\widetilde{\psi}(u)$ because of the choice of the quadratic splines. We will later describe the choice of K_n, M_n, and the knots for defining B-splines.

4.2.3 Computational Algorithm

We design an algorithm that maximizes (4.2) by casting the transformation model into the framework of missing values. Then we can immediately apply the EM algorithm to facilitate the computation. Within each loop of the EM algorithm, the optimization problem reduces to the single-index weighted proportional hazards model, which can be solved by

an interative procedure similar to that in Chapter 3. The detailed algorithm is describe as follows.

In the initial step, we fit a Cox model with covariates $(\boldsymbol{X}^{\mathrm{T}}, A, A\boldsymbol{Z}^{\mathrm{T}})^{\mathrm{T}}$, and initialize $\widehat{\boldsymbol{\alpha}}_0$ and $\widehat{\boldsymbol{\beta}}_0$ with the estimated coefficient for $\boldsymbol{X}$ and the normalized coefficient for $A\boldsymbol{Z}$, respectively. The initial value for $\boldsymbol{\gamma}$ is set to be the least square estimation coefficients in the B-spline approximation to a linear function, where the intercept and slope of this function are the coefficient of A and the Euclidean norm of the coefficient for $A\boldsymbol{Z}$ in the Cox model, respectively. We initialize the cumulative hazard $\widehat{\Lambda}_0$ to be the Breslow estimator.

At the lth iteration, using the current parameter values $\widehat{\boldsymbol{\alpha}}_{l-1}, \widehat{\boldsymbol{\beta}}_{l-1}$, $\widehat{\boldsymbol{\gamma}}_{l-1}$, $\widehat{\Lambda}_{l-1}$ and $\widehat{r}_{l-1}$, we compute $\{W_{l-1,i}\}$ and set the link function as $\widehat{\psi}_{l-1}(u) \equiv \widehat{\boldsymbol{\gamma}}_{l-1}^{\mathrm{T}}\boldsymbol{B}(u)$. We first update $\boldsymbol{\beta}$ by maximizing

$$\sum_{i=1}^{n}\Delta_i\Bigg\{\widehat{\boldsymbol{\alpha}}_{l-1}^{\mathrm{T}}\boldsymbol{X}_i + A_i\{\widehat{\psi}_{l-1}(\widehat{\boldsymbol{\beta}}_{l-1}^{\mathrm{T}}\boldsymbol{Z}_i) + \widehat{\psi}'_{l-1}(\widehat{\boldsymbol{\beta}}_{l-1}^{\mathrm{T}}\boldsymbol{Z}_i)(\boldsymbol{\beta}-\widehat{\boldsymbol{\beta}}_{l-1})^{\mathrm{T}}\boldsymbol{Z}_i\}-$$
$$\log\Bigg(\sum_{j=1}^{n}Y_j(\widetilde{T}_i)W_{l-1,j}\exp\Big[\widehat{\boldsymbol{\alpha}}_{l-1}^{\mathrm{T}}\boldsymbol{X}_j + A_j\{\widehat{\psi}_{l-1}(\widehat{\boldsymbol{\beta}}_{l-1}^{\mathrm{T}}\boldsymbol{Z}_j) + \widehat{\psi}'_{l-1}(\widehat{\boldsymbol{\beta}}_{l-1}^{\mathrm{T}}\boldsymbol{Z}_j)(\boldsymbol{\beta}-\widehat{\boldsymbol{\beta}}_{l-1})^{\mathrm{T}}\boldsymbol{Z}_j\}\Big]\Bigg)\Bigg\},$$

subject to the constraint $\|\boldsymbol{\beta}\| = 1$. Essentially, we approximate $\widehat{\psi}_{l-1}(\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z})$ in the partial likelihood function 4.5 by the first order Taylor expansion at $\widehat{\boldsymbol{\beta}}_{l-1}$, which results in a concave function of $\boldsymbol{\beta}$. This optimization step is solved through the Lagrange multiplier method. We check that the partial likelihood increases at $\widehat{\boldsymbol{\beta}}_l$ compared to $\widehat{\boldsymbol{\beta}}_{l-1}$. Otherwise, we calculate the partial likelihood at $2^{-q}\widehat{\boldsymbol{\beta}}_l + (1-2^{-q})\widehat{\boldsymbol{\beta}}_{l-1}$, for $q = 1, 2, \cdots$, and update $\boldsymbol{\beta}$ as $2^{-q}\widehat{\boldsymbol{\beta}}_l + (1-2^{-q})\widehat{\boldsymbol{\beta}}_{l-1}$ with the smallest q that increases the partial likelihood. Next, given $\widehat{\boldsymbol{\beta}}_l$, we update $\boldsymbol{\alpha}$ and $\boldsymbol{\gamma}$ by maximizing $L_3(\boldsymbol{\alpha}, \widehat{\boldsymbol{\beta}}_l, \boldsymbol{\gamma})$ under the constraint $-M_n \leqslant \gamma_1 \leqslant \cdots \leqslant \gamma_{K_n+1} \leqslant M_n$. We note that the object functions in this optimization is strictly concave and that the constraint sets are convex. Therefore, there are guaranteed unique global maximums and many packages for convex optimization can be used. Particularly, we adopt R package

quadprog in our numerical studies. Last, we update $\widehat{\Lambda}$ by

$$\widehat{\Lambda}_l(t) = \sum_{i=1}^{n} \frac{I(\widetilde{T}_i \leqslant t)\Delta_i}{\sum_{j=1}^{n} Y_j(\widetilde{T}_i) W_{l-1,j} \exp\{\widehat{\boldsymbol{\alpha}}_l^{\mathrm{T}} \boldsymbol{X}_j + A_j \widehat{\psi}(\widehat{\boldsymbol{\beta}}_l^{\mathrm{T}} \boldsymbol{Z}_j)\}}.$$

We iterate until convergence. The algorithm is guaranteed to converge, since the objective function 4.4 in the maximization step is increased at each iteration, and will only be unchanged upon convergence. Let $\widehat{\boldsymbol{\alpha}}$, $\widehat{\boldsymbol{\beta}}$, $\widehat{\boldsymbol{\gamma}}$, and $\widehat{\Lambda}$ be the final estimates of $\boldsymbol{\alpha}$, $\boldsymbol{\beta}$, $\boldsymbol{\gamma}$, and Λ, respectively. The link function $\psi(u)$ is estimated by $\widehat{\psi}(u) \equiv \widehat{\boldsymbol{\gamma}}^{\mathrm{T}} \boldsymbol{B}(u)$. Finally, the optimal treatment rule is estimated as $I\{\widehat{\boldsymbol{\beta}}^{\mathrm{T}} \boldsymbol{Z} < \widehat{\psi}^{-1}(0)\}$.

4.2.4 Variance estimation

We use profile likelihood (Murphy & van der Vaart, 2000) to estimate the covariance matrix of $(\widehat{\boldsymbol{\alpha}}^{\mathrm{T}}, \widehat{\boldsymbol{\beta}}_{-q}^{\mathrm{T}})^{\mathrm{T}}$, which we denote as $\boldsymbol{\theta}$ thereafter. Specifically, we define the profile loglikelihood

$$pl_n(\boldsymbol{\theta}) = \max_{\Lambda \in \mathcal{C}, \psi \in \mathcal{D}} logL_n(\boldsymbol{\theta}, \psi, \Lambda),$$

where $\mathcal{C}$ is the set of step functions with nonnegative jumps at $\widetilde{T}_i$ with $\Delta_1 = 1$, and $\mathcal{D}$ is the set of monotone increasing functions spanned by the B-splines. Then the covariance matrix of $\boldsymbol{\theta}$ is estimated by the negative inverse of the matrix whose (j,k)th element is

$$h_n^{-2} \left\{ pl_n(\widehat{\boldsymbol{\theta}}) - pl_n(\widehat{\boldsymbol{\theta}} + h_n \boldsymbol{e}_k) - pl_n(\widehat{\boldsymbol{\theta}} + h_n \boldsymbol{e}_j) + pl_n(\widehat{\boldsymbol{\theta}} + h_n \boldsymbol{e}_k + h_n \boldsymbol{e}_j) \right\}, \tag{4.6}$$

where $\boldsymbol{e}_j$ is the jth canonical vector in $\mathbb{R}^{p+q-1}$ and h_n is a constant of order $n^{-1/2}$. To calculate $pl_n(\boldsymbol{\theta})$, we reuse the proposed EM algorithm with $\boldsymbol{\theta}$ held fixed and only updating $\widehat{\Lambda}$ and $\widehat{\psi}$ in the M-step.

Numerical results show that the second order perturbation formula might lead to negative results in rare cases, and the rate $n^{-1/2}$ might not work the best in small sample scenarios. Therefore, we also propose the first order analog to the above formula, with h_n of order n^{-1}. Specifically, the covariance matrix of $\boldsymbol{\theta}$ is estimated by the inverse of the matrix whose

(j,k)th element is

$$h_n^{-2}\sum_{i=1}^{n}\left\{pl_n^i(\widehat{\boldsymbol{\theta}})-pl_n^i(\widehat{\boldsymbol{\theta}}+h_n\boldsymbol{e}_k)\right\}\left\{pl_n^i(\widehat{\boldsymbol{\theta}})-pl_n^i(\widehat{\boldsymbol{\theta}}+h_n\boldsymbol{e}_j)\right\}, \tag{4.7}$$

where pl_n^i is the contribution of the ith subject to the profile likelihood.

4.3 Asymptotic Properties

Let $\boldsymbol{\alpha}_0 \equiv (\alpha_{0,1},\cdots,\alpha_{0,p})^{\mathrm{T}}$, $\boldsymbol{\beta}_0 \equiv (\beta_{0,1},\cdots,\beta_{0,q})^{\mathrm{T}}$, and $\psi_0(\cdot)$ denote the true values of $\boldsymbol{\alpha},\boldsymbol{\beta}$, and $\psi(\cdot)$, respectively. In addition, let τ denote the study duration and $\mathcal{Z}$ denote the union of the support of $\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z}$ for all $\|\boldsymbol{\beta}\|=1$. We impose the following regularity conditions.

Condition 1. The true values $\boldsymbol{\alpha}_0$ and $\boldsymbol{\beta}_0$ lie in the interior of a known compact set; $\psi_0(u)$ is a strictly increasing function of u and three-times differentiable in $\mathcal{Z}$, and $\lambda_0(t)$ is twice differentiable in $[0,\tau]$.

Condition 2. The conditional density of C given $\boldsymbol{X}$ has bounded twice differentiable derivative on its support $[0,\tau]$. The conditional distribution of $(A,\boldsymbol{X})$ given $\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z}$ has a continuously differentiable density with respect to a dominating measure.

Condition 3. With probability 1, $P(Y(\tau)=1|\boldsymbol{X})>c_0$ for some positive constant c_0. In addition, if $\boldsymbol{c}_1^{\mathrm{T}}\boldsymbol{X}=d_1$ with probability 1 for some constant vector $\boldsymbol{c}_1$ and constant d_1, then $\boldsymbol{c}_1=\boldsymbol{0}$ and $d_1=0$.

Condition 4. The number of the knots satisfies that $K_n\to\infty$ and $n^{-1/2}K_n^7\to\infty$. In addition, the adjacent distance of the interior knots is between $c^{-1}K_n^{-1}$ and cK_n^{-1} for some positive constant c.

Condition 5. For any positive c_0, $\limsup_{x\to\infty}\{G(c_0x)\}^{-1}\log\{x\sup_{y\leqslant x}G'(y)\}\leqslant 1$.

Remark 1. The second part of Condition 1 and Condition 2 ensure smoothness for the functions ψ_0 and λ_0, the conditional distribution of C given $\boldsymbol{X}$, and the conditional distribution of $(A,\boldsymbol{X})$ given $\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z}$. Condition 3 ensures that a non-trivial proportion of subjects is censored at τ and that $(1,\boldsymbol{X}^{\mathrm{T}})^{\mathrm{T}}$ is linearly independent.

Remark 2. Condition 4 is satisfied by most transformations. For example, the limit goes to zero for the Box-Cox transformation $G(x) = \{(1+x)^\rho - 1\}/\rho$ with $\rho > 0$ and goes to one for the logarithmic transformation $G(x) = r^{-1}\log(1+rx)$.

We state the consistency and asymptotic distribution of the estimators for the model parameters in the following three theorems, whose proofs are given in the Appendix.

Theorem 4.3.1. *Under Conditions 1–4,* $||\widehat{\psi} - \psi_0||_{W^{1,\infty}(\mathcal{Z})} + \sup_{t\in[0,\tau]} |\widehat{\Lambda} - \Lambda_0| + ||\widehat{\boldsymbol{\alpha}} - \boldsymbol{\alpha}_0|| + ||\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0|| \to 0$ *in probability, where for any differentiable function* f *with derivative* f', $||f||_{W^{1,\infty}(\mathcal{Z})}$ *is defined as* $||f||_{L^\infty(\mathcal{Z})} + ||f'||_{L^\infty(\mathcal{Z})}$. *Furthermore,* $||\widehat{\boldsymbol{\alpha}} - \boldsymbol{\alpha}_0||^2 + ||\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0||^2 + ||\widehat{\psi} - \psi_0||^2_{L_2(\mathcal{Z})} = o_p(n^{-1/2})$.

To describe the asymptotic distribution, we assume $\beta_{0q} > 0$ without loss of generality. For a q-dimensional vector $\boldsymbol{x}$, let $\boldsymbol{x}_{-q} = (x_1, \cdots, x_{q-1})^{\mathrm{T}}$.

Theorem 4.3.2. *Under Conditions 1–5,* $n^{1/2}(\widehat{\boldsymbol{\alpha}}^{\mathrm{T}} - \boldsymbol{\alpha}_0^{\mathrm{T}}, \widehat{\boldsymbol{\beta}}_{-q}^{\mathrm{T}} - \boldsymbol{\beta}_{0,-q}^{\mathrm{T}})^{\mathrm{T}}$ *converges in distribution to a zero-mean normal random vector with covariance matrix* $\boldsymbol{\Sigma}^{-1}$, *and* $\boldsymbol{\Sigma}^{-1}$ *achieves the semiparametric efficiency bound.*

4.4 Simulation Studies

We conducted simulation studies to assess the performance of the proposed methods under realistic scenarios. We considered sample sizes $n = 500$, 1000, and 2000. The baseline hazard function $\Lambda(t)$ was assumed to follow a Weibull distribution with shape parameter 2.5 and scale parameter 2. We considered two link functions: (a) (exponential) $\psi_0(u) = e^u - 0.5$; (b) (linear) $\psi_0(u) = u + 0.4$; and (c) (sine)

$$\psi_0(u) = \begin{cases} -3 \text{ if } u \leqslant -1/2 \\ \sin(\pi u) - 1 \text{ if } -1/2 < u \leqslant 1/2 \\ 1 \text{ if } u > 1/2. \end{cases}$$

The censoring time was assumed to follow Unif$[0, \tau]$, where τ was chosen such that the censoring rate was 50%. We generated four independent covariates from Unif$[-1, 1]$ for the main effects and assumed that the treatment effect on the survival time depends on the first two covariates with the main effects. We set the corresponding main effect parameters $(\alpha_1, \alpha_2, \alpha_3, \alpha_4)^{\mathrm{T}}$ to $(-0.4, -0.2, 0.2, 0.4)^{\mathrm{T}}$ and the interaction effect parameters $(\beta_1, \beta_2, \beta_3, \beta_4)^{\mathrm{T}}$ to $(-0.5, -0.5, 0.5, 0.5)^{\mathrm{T}}$. The treatment assignment A was assumed to be independent of $\boldsymbol{X}, \boldsymbol{Z}$ and to follow a Bernoulli distribution with success probability of 0.5. We considered a series of models with the number of knots K_n ranging from 3 to 9. The majority of the 10,000 replicates selected a final model with $K_n = 3, 3$, and 6 by the AIC criterion for the exponential, linear, and sine link functions, respectively. In the variance calculation, we used $h_n = n^{-1}$ as the perturbation parameter.

Table 4.1-4.2 summarizes the results for the estimation of $(\beta_1, \beta_2, \beta_3, \beta_4)^{\mathrm{T}}$ under the proposed model with different links and transformation models. Note that the scenario $r = 0$ and $r = 1$ corresponds to the proportional hazards single-index model and proportional odds single-index model, respectively. The biases of the parameter estimators under our model are small and decrease as the sample size increases increases. The variance estimators and the corresponding confidence intervals for $\boldsymbol{\beta}$ become more accurate as the sample size increases to 2000 or higher.

4.5 Application to evaluate adjuvant therapy of large-bowel carcinoma

The data is from a randomized clinical trial to evaluate the surgical adjuvant therapy of large-bowel carcinoma (Moertel et al., 1990; Laurie et al., 1989). This was a randomized clinical trial on patients with resected stages B and C colorectal carcinoma. The main goal was to assess the effectiveness of fluorouracil (5-FU) plus levamisole versus levamisole alone versus the observation group. The time to death and disease relapse were recoreded, with a median follow-up time of 6.5 years. 50% of the subjects were censored.

We aimed to investigate the interaction between baseline covariates and treatment effect on the time to disease relapse and derive the optimal treatment regimen for each patient

in the study. With some exploratory analysis, we included the following 8 covariates for the main effect: location of primary neoplasm, extent of local speed, perforation, adherence, number of positive nodes, age, number of days from tumor resection to start of treatment and pre-operative CEA level. Among these, perforation, adherence, number of positive nodes, age and pre-operative CEA level were included for the interaction terms. In the exploratory analysis, we found that the performance in the levamisole alone group and the observation group is comparable. Therefore, we combined those patients and coded them as receiving treatment $A = 0$, and the treatment of fluorouracil (5-FU) plus levamisole was labeled as $A = 1$.

We fitted the transformation single index model with the logarithmic transformation parameter ranging from 0 to 2, with a step size of 0.1. We evaluated the the performance of different single-index transformation models using a five-fold cross validation. Each time, we fitted the model with the left-out dataset, and computed two performance benchmarks with the left-out: the C-index and the p-value from the log-rank test. Each time in the cross validation, we calculated the optimal treatment assignment for each patient in the left-out dataset. The C-index measured if the risk of the disease relapse as predicted in the model is concordant with their actual order of timing for every pair of patients in the left-out dataset. The p-value in the log-rank test was obtained in testing the difference of disease relapse over time between patients who followed the treatment versus those who did not follow the treatment. The C-index benchmark generally favored a larger r. The log rank test suggested a model with r in the range of 0.8 and 1, and r greater than 1 was almost as good as the optimal r. Combining these two results, we decided to choose $r = 1.5$ as our final transformation parameter. Details of the cross validation criteria is provided in Figure 4.1.

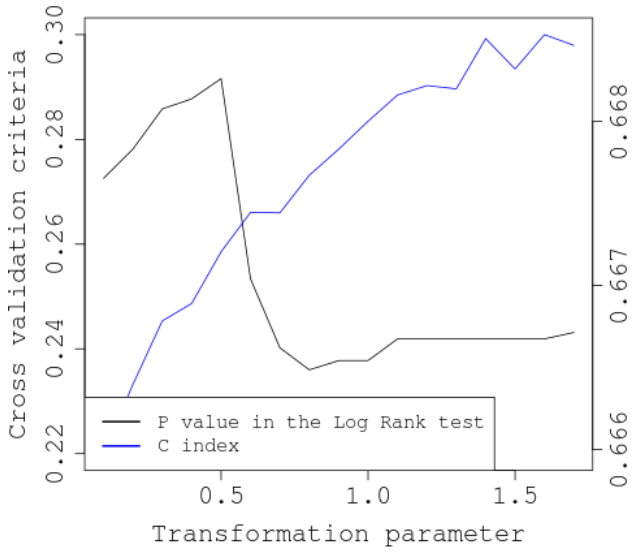

Figure 4.1: Cross validation criteria in choosing the best transformation.

Table 4.3 displays the estimation results for the treatment-covariate interaction effects under the proposed model and the single-index proportional hazards model. Both models indicate that age is significantly related to the treatment effect on time to relapse, with older patients more likely to benefit from fluorouracil (5-FU) plus levamisole than levamisole alone. Perforation was a deterministic factor in the optimal treatment rule estimation under the single-index model, but did not show the effect with the proposed model. Adherence was marginally significant under the proposed model but not so under the single-index model. A comparison of the estimated link functions is given in Figure 4.2.

In order to measure the performance of the proposed model in assigning patients to the optimal treatment rule, we compared the probability of disease relapse over time for patient who followed the treatment rule versus those who did not follow the treatment rule. It is clear in Figure 4.3 that the patients have a much smaller risk of experiencing disease relapse if they follow the recommended treatment strategy. That is, the treatment rule estimated by the proposed transformation single-index model benefits the patients in delaying their time to experience the disease relapse.

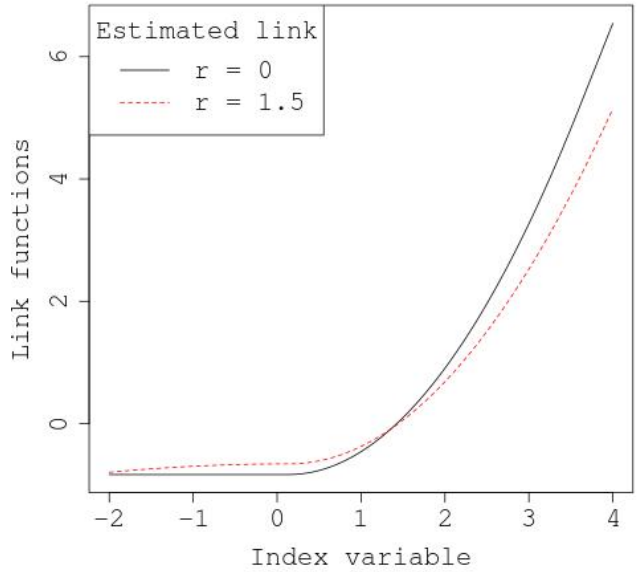

Figure 4.2: Estimation of the link functions under the transformation single-index model ($r = 1.5$) and the single-index proportional hazards model ($r = 0$).

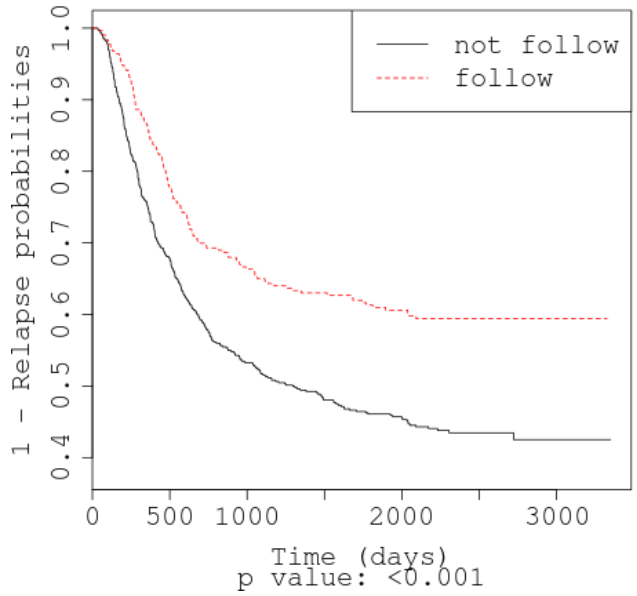

Figure 4.3: Estimation of the survival curve for patients who follow the treatment recommendation under the proposed model and those who do not follow the treatment recommendation.

4.6 Discussion

In practice, the transformation function is an unknown piece in the model and needs to be selected. In the real data application, we used p value in the log rank test and the C-index as benchmarks. These statistics are estimated by the cross validation approach. Other criteria, such as the log-likelihood, may also work. However, with the aim of estimating the optimal treatment rule, we think the criteria that emphasize the treatment assignment is more relevant.

Although this chapter is mainly focused on estimation of the interaction parameters and the optimal treatment rule under the transformation model, the proposed method can be easily extend for inference of the value function for each patient. Asymptotic properties for a predefined loss or benefit function that measures how much benefit an individual receives from the treatment can be obtained using strategies similar to the proof for the asymptotic behavior of the interaction parameters.

One practical issue in using the proposed model is to determine which variables to include for the potential interaction term. In the next chapter, we will develop a variable selection approach that allows high dimensional input and automatically differentiates the important and unimportant variables in predicting the optimal treatment rule.

Table 4.1: Simulation results for the estimation of $\boldsymbol{\beta} = (-0.5, -0.5, 0.5, 0.5)$ with different link functions.

			$r = 0$				$r = 1$			
Link	**n**		**Bias**	**SE**	**SEE**	**CP**	**Bias**	**SE**	**SEE**	**CP**
Exponential	1000	β_1	0.02	0.10	0.09	0.90	0.03	0.14	0.12	0.89
		β_2	0.01	0.09	0.09	0.91	0.02	0.14	0.12	0.90
		β_3	0.00	0.09	0.09	0.93	-0.02	0.14	0.12	0.91
		β_4	-0.01	0.09	0.09	0.92	-0.02	0.14	0.13	0.90
	2000	β_1	0.01	0.07	0.06	0.93	0.01	0.09	0.09	0.93
		β_2	0.01	0.06	0.06	0.93	0.02	0.10	0.09	0.92
		β_3	0.00	0.06	0.06	0.94	-0.01	0.09	0.09	0.93
		β_4	0.00	0.06	0.06	0.95	0.00	0.09	0.09	0.94
	5000	β_1	0.00	0.04	0.04	0.95	0.00	0.06	0.06	0.96
		β_2	0.00	0.04	0.04	0.95	0.01	0.06	0.06	0.96
		β_3	0.00	0.04	0.04	0.95	-0.01	0.06	0.06	0.95
		β_4	0.00	0.04	0.04	0.94	0.00	0.06	0.06	0.94
Linear	1000	β_1	0.03	0.15	0.12	0.87	0.05	0.21	0.17	0.85
		β_2	0.02	0.15	0.12	0.87	0.05	0.21	0.17	0.86
		β_3	-0.02	0.15	0.12	0.88	-0.04	0.21	0.17	0.86
		β_4	-0.02	0.15	0.15	0.88	-0.05	0.21	0.19	0.88
	2000	β_1	0.01	0.10	0.09	0.93	0.02	0.14	0.13	0.93
		β_2	0.01	0.10	0.09	0.92	0.03	0.14	0.13	0.91
		β_3	-0.01	0.10	0.09	0.93	-0.02	0.14	0.13	0.93
		β_4	-0.01	0.10	0.09	0.93	-0.01	0.14	0.13	0.93
	5000	β_1	0.01	0.06	0.06	0.96	0.01	0.08	0.09	0.97
		β_2	0.00	0.06	0.06	0.94	0.01	0.09	0.09	0.95
		β_3	-0.01	0.06	0.06	0.95	-0.02	0.09	0.09	0.96
		β_4	0.00	0.06	0.06	0.93	0.00	0.09	0.09	0.94

r denotes the transformation parameter in the logarithmic transformation. Bias and SE are the bias and standard error of the parameter estimator, respectively; SEE is the mean of the standard error estimator; CP is the coverage probability of the 95% confidence interval.

Table 4.2: Simulation results for the estimation of $\boldsymbol{\beta} = (-0.5, -0.5, 0.5, 0.5)$ with different link functions.

			$\boldsymbol{r} = \mathbf{0}$				$\boldsymbol{r} = \mathbf{1}$			
Link	**n**		**Bias**	**SE**	**SEE**	**CP**	**Bias**	**SE**	**SEE**	**CP**
Sine	1000	β_1	0.01	0.07	0.06	0.93	0.02	0.13	0.10	0.87
		β_2	0.00	0.07	0.06	0.93	0.01	0.12	0.09	0.88
		β_3	0.00	0.07	0.06	0.93	-0.01	0.13	0.09	0.87
		β_4	-0.01	0.07	0.06	0.92	-0.02	0.13	0.10	0.86
	2000	β_1	0.00	0.04	0.05	0.96	0.01	0.07	0.07	0.94
		β_2	0.00	0.04	0.05	0.95	0.01	0.08	0.07	0.93
		β_3	0.00	0.05	0.05	0.94	0.00	0.07	0.07	0.93
		β_4	0.00	0.05	0.05	0.93	-0.01	0.08	0.07	0.92
	5000	β_1	0.00	0.03	0.03	0.96	0.00	0.04	0.05	0.96
		β_2	0.00	0.03	0.03	0.96	0.00	0.04	0.05	0.96
		β_3	0.00	0.03	0.03	0.95	0.00	0.04	0.05	0.95
		β_4	0.00	0.03	0.03	0.95	0.00	0.04	0.05	0.96

See note to Table 4.1.

Table 4.3: Estimation of the parameters under different transformation models

	$r = 0$			$r = 1.5$		
Parameters	**Est.**	**Std. Err.**	***P*-value**	**Est.**	**Std. Err.**	***P*-value**
Main Effects						
Sigmoid neoplasm	-0.18	0.11	0.09	-0.26	0.17	0.11
Cecum neoplasm	-0.14	0.12	0.27	-0.17	0.19	0.36
Muscular Spread	0.31	0.53	0.56	0.23	0.61	0.71
Serosa Spread	0.85	0.50	0.09	0.92	0.58	0.11
Contuiguos structures	1.44	0.53	0.01	1.82	0.67	0.01
Perforation	0.14	0.26	0.59	0.12	0.45	0.79
Adherence	0.18	0.13	0.18	0.20	0.23	0.39
Log positive nodes	0.39	0.05	< 0.01	0.54	0.07	< 0.01
Age	0.01	0.05	0.82	0.02	0.08	0.79
Days to treatment	-0.08	0.05	0.08	-0.08	0.07	0.26
Pre-operative CEA level	0.09	0.05	0.08	0.14	0.08	0.07
Interactions						
Perforation	0.25	1.24	0.84	0.04	1.63	0.98
Adherence	0.56	0.45	0.21	0.65	0.36	0.07
Log positive nodes	0.28	0.18	0.12	0.24	0.23	0.30
Age	-0.67	0.19	< 0.01	-0.67	0.30	0.03
Pre-operative CEA level	0.32	0.17	0.06	0.28	0.25	0.26

The transformation parameter $r = 0$ corresponds to a single-index proportional hazards model. “Est." and "Std. Err." denote the estimate and the standard error estimate for a parameter.

CHAPTER 5: VARIABLE SELECTION WITH SINGLE-INDEX MODELS WITH APPLICATION TO OPTIMAL TREATMENT REGIMES

5.1 Introduction

With the modern development of technology, wide use of electronic devices, and the availability of electronic health record, researchers are able to collect more variables from each patient than any time in the past. The detailed information provides both the opportunity and the challenge. It is an opportunity for researchers to discover unknown associations between treatment benefit and the patient characteristics and thus develop a more effective treatment procedure. It is challenging because among the massive information, only a small portion is truly important in predicting the optimal treatment rule. Such variables need to be identified from the noise in a computational efficient way with a sample size in a typical clinical trial, which is usually not very large.

This procedure is usually termed as variable selection or dimension reduction and has been studied using stepwise model selection (Chatfield, 1995; Harrell et al., 1996; Steyerberg et al., 1999), the Akaike information criterion (Akaike, 1970), and the Bayesian information criterion (Schwarz et al., 1978). Certain penalties, such as the adaptive lasso (Zou, 2006) and the smoothly clipped absolute deviation (Fan and Li, 2001), have been shown to enjoy the oracle penalty when they are used for variable selection with non-censored outcomes. For censored outcomes, variable selection with the adaptive lasso penalty has been discussed in the Zhang and Lu (2007) and Liu and Zeng (2013), among other methods.

In the field of personalized medicine, there has been some discussion of variable selection with non-censored outcomes. Fan et al. (2016) studied the penalized A-learning in a multiple-stage treatment where a Dantzig selector (Candes et al., 2007) directly penalizes the A-leaning estimating equations. Gunter et al. (2011) discussed variable selection for qualitative

interactions using the Adjusted Gain in Value (AGV) lasso penalty. Nezhad et al. (2016) designed the integrated feature selection method with stacked autoencoders in deep learning to achieve dimension reduction. All of the above mentioned works can only handle non-censored outcomes. To our knowledge, there is no existing variable selection method in estimating the personalized treatment rule with censored outcomes.

In this chapter, we propose a variable selection approach by adding the adaptive lasso penalty to the log-likelihood of a single-index transformation model. This model encompasses both the single-index proportional hazards model and the single-index proportional odds model as special cases, among many other models in between. We design an estimation procedure by casting the problem into a weighted Cox model framework. We show the oracle penalty and asymptotic distribution of the adaptive lasso estimator. Performance of the proposed method is demonstrated by simulation examples and a clinical trial application.

5.2 Methods

5.2.1 Data and Model

Let T denote the survival time, A denote the binary treatment with values 0 and 1, and $\boldsymbol{X}$ denote the baseline covariates assumed to be bounded. To characterize the heterogeneity of treatment effects under the semiparametric transformation model, we assume that the cumulative hazard function for T conditional on A and $\boldsymbol{X}$ takes the form

$$\Lambda(t; A, \boldsymbol{X}) = G\left[\int_0^t \exp\{\boldsymbol{\alpha}^{\mathrm{T}}\boldsymbol{X} + A\psi(\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z})\}d\Lambda(s)\right], \tag{5.1}$$

where $G(\cdot)$ is a specific transformation function that is strictly increasing, $\Lambda(\cdot)$ is an unknown increasing function (Zeng & Lin, 2006), $\boldsymbol{Z}$ is a subset of $\boldsymbol{X}$, $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ are the regression parameters, and $\psi(\cdot)$ is an unknown and strictly increasing link function. For identifiability, we further assume that there are non-zero components in $\boldsymbol{\beta}$ and $||\boldsymbol{\beta}|| = 1$, where $||\boldsymbol{x}||$ is the Euclidean norm of vector $\boldsymbol{x}$.

Clearly, when $\psi(x) = x$, model (5.1) reduces to the usual transformation model with the

interaction between the treatment and the covariates. The choice of $G(x) = x$ yields the proportional hazards models with a single index term for treatment covariate interaction. An important class of transformation models is the class of logarithmic transformations $G(x) = r^{-1}\log(1 + rx)(r \geqslant 0)$, with $r = 0$ and $r = 1$ corresponding to the proportional hazards and proportional odds models, respectively. Since all the covariates are time-independent, model (5.1) can be rewritten as a linear transformation model $\log \Lambda(T) = -\boldsymbol{\alpha}^{\mathrm{T}}\boldsymbol{X} - A\psi(\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z}) + \epsilon$, where ϵ is an error term with distribution function $1 - \exp[-G\{\exp(x)\}]$ (Chen et al., 2002).

5.2.2 Initial Estimate

Let $(\widetilde{T}_i, \Delta_i, A_i, \boldsymbol{X}_i)$ $(i = 1, \ldots, n)$ denote the data from a randomized trial consist of n right-censored observations, where $\widetilde{T}_i = \min(T_i, C_i)$ and $\Delta_i = I(T_i \leqslant C_i)$, and C_i is the potential censoring time for subject i. We make the standard independent censoring assumption that C_i is independent of T_i given $(A_i, \boldsymbol{X}_i)$. The observed log-likelihood function concerning the model parameters is given by

$$\sum_{i=1}^{n} \Big(\Delta_i \left\{\boldsymbol{\alpha}^{\mathrm{T}}\boldsymbol{X}_i + A_i\psi(\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z}_i)\right\} + \Delta_i \log \lambda(\widetilde{T}_i) + \Delta_i \log G'\left[\exp\{\boldsymbol{\alpha}^{\mathrm{T}}\boldsymbol{X}_i + A_i\psi(\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z}_i)\}\Lambda(\widetilde{T}_i)\right] - G\left[\exp\{\boldsymbol{\alpha}^{\mathrm{T}}\boldsymbol{X}_i + A_i\psi(\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z}_i)\}\Lambda(\widetilde{T}_i)\right]\Big), \quad (5.2)$$

where $\Lambda(t) = \int_0^t \lambda(s)ds$.

Expression (5.2) involves both $\boldsymbol{\alpha}, \boldsymbol{\beta}$ and the infinite dimensional parameters ψ and Λ, and it may not be concave in these parameters. Also, there is no partial likelihood function available due to the transformation G. As shown in Chapter 4, the non-parametric maximum likelihood estimators can be obtained using a combination of the expectation-maximization (EM) algorithm and sieve estimation approach. In each iteration in the EM algorithm, the posterior expectation in the E step can be written as

$$W_i = G'\left\{\exp\{\widetilde{\boldsymbol{\alpha}}^{\mathrm{T}}\boldsymbol{X}_i + A_i\widetilde{\psi}(\widetilde{\boldsymbol{\beta}}^{\mathrm{T}}\boldsymbol{Z}_i)\}\widetilde{\Lambda}(\widetilde{T}_i)\right\} - \Delta_i \frac{G''\left\{\exp\{\widetilde{\boldsymbol{\alpha}}^{\mathrm{T}}\boldsymbol{X}_i + A_i\widetilde{\psi}(\widetilde{\boldsymbol{\beta}}^{\mathrm{T}}\boldsymbol{Z}_i)\}\widetilde{\Lambda}(\widetilde{T}_i)\right\}}{G'\left\{\exp\{\widetilde{\boldsymbol{\alpha}}^{\mathrm{T}}\boldsymbol{X}_i + A_i\widetilde{\psi}(\widetilde{\boldsymbol{\beta}}^{\mathrm{T}}\boldsymbol{Z}_i)\}\widetilde{\Lambda}(\widetilde{T}_i)\right\}},$$

where $\Lambda(t)$ denotes the step size of Λ at t. The function we aim to maximize in the M step takes the form

$$L(\boldsymbol{\alpha}, \boldsymbol{\beta}, \psi) = \sum_{i=1}^{n} \Delta_i \left(\boldsymbol{\alpha}^{\mathrm{T}} \boldsymbol{X}_i + A_i \boldsymbol{\gamma}^{\mathrm{T}} \boldsymbol{B}(\boldsymbol{\beta}^{\mathrm{T}} \boldsymbol{Z}_i) - \log \left[\sum_{j=1}^{n} Y_j(\widetilde{T}_i) W_j \exp\{\boldsymbol{\alpha}^{\mathrm{T}} \boldsymbol{X}_j + A_j \boldsymbol{\gamma}^{\mathrm{T}} \boldsymbol{B}(\boldsymbol{\beta}^{\mathrm{T}} \boldsymbol{Z}_j)\} \right] \right). \tag{5.3}$$

where $\gamma_1 \leqslant \cdots \leqslant \gamma_{K_n+1}$, $\|\boldsymbol{\beta}\| = 1$, and $\boldsymbol{B}(u) = (B_1(u), \cdots, B_{K_n+1}(u))^{\mathrm{T}}$ denote quadratic B-spline bases corresponding to K_n distinct knots in an interval containing the support of $\boldsymbol{\beta}^{\mathrm{T}} \boldsymbol{Z}$ for any $\|\boldsymbol{\beta}\| = 1$.

5.2.3 Variable Selection

We consider the adaptive lasso penalty for variable selection for its oracle properties. To accommodate penalties into the original optimization problem, we first examine the equation (5.3), which resembles the form of the partial log-likelihood function for a Cox model. Therefore, we propose to adapt the ideas in Liu and Zeng (2013) and Zhang and Lu (2007) in the estimation with the adaptive lasso penalty. Specifically, after we apply the penalty, the objective function can be written as

$$Q_n(\boldsymbol{\beta}) = -\sum_{i=1}^{n} \Delta_i \left(\boldsymbol{\alpha}^{\mathrm{T}} \boldsymbol{X}_i + A_i \psi(\boldsymbol{\beta}^{\mathrm{T}} \boldsymbol{Z}_i) - \log \left[\sum_{j=1}^{n} Y_j(\widetilde{T}_i) W_j \exp\{\boldsymbol{\alpha}^{\mathrm{T}} \boldsymbol{X}_j + A_j \psi(\boldsymbol{\beta}^{\mathrm{T}} \boldsymbol{Z}_j)\} \right] \right) + n\rho \sum_{j=1}^{q} |\widetilde{\beta}_j|^{-1} |\beta_j|, \tag{5.4}$$

subject to $||\boldsymbol{\beta}|| = 1$, where ρ is a tuning parameter and $\widetilde{\boldsymbol{\beta}}$ is the initial estimate. A large value of the initial estimate indicates higher predictive power and the corresponding variable receives less penalty; a smaller initial estimate implies less importance of the variable in the model fitting and is penalized with a higher weight.

We approximate $\psi(\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z})$ by the first order Taylor expansion at $\widetilde{\boldsymbol{\beta}}$ in equation (5.3):

$$
\begin{aligned}
l_n(\boldsymbol{\beta}) = \sum_{i=1}^{n} \Delta_i \Big\{ & \boldsymbol{\alpha}^{\mathrm{T}}\boldsymbol{X}_i + A_i\{\psi(\widetilde{\boldsymbol{\beta}}^{\mathrm{T}}\boldsymbol{Z}_i) + \psi(\widetilde{\boldsymbol{\beta}}^{\mathrm{T}}\boldsymbol{Z}_i)(\boldsymbol{\beta} - \widetilde{\boldsymbol{\beta}})^{\mathrm{T}}\boldsymbol{Z}_i\} - \\
& \log\Big(\sum_{j=1}^{n} Y_j(\widetilde{T}_i) W_j \exp\Big[\boldsymbol{\alpha}^{\mathrm{T}}\boldsymbol{X}_j + A_j\{\psi(\widetilde{\boldsymbol{\beta}}^{\mathrm{T}}\boldsymbol{Z}_j) + \psi'(\widetilde{\boldsymbol{\beta}}^{\mathrm{T}}\boldsymbol{Z}_j)(\boldsymbol{\beta} - \widetilde{\boldsymbol{\beta}})^{\mathrm{T}}\boldsymbol{Z}_j\}\Big]\Big)\Big\}.
\end{aligned}
$$

In terms of $\boldsymbol{\beta}$, this expression takes the form of a weighted partial log-likelihood function with the ith observation's weight given by $W_j \exp[\boldsymbol{\alpha}^{\mathrm{T}}\boldsymbol{X}_j + A_j\{\psi(\widetilde{\boldsymbol{\beta}}^{\mathrm{T}}\boldsymbol{Z}_j) - \psi'(\widetilde{\boldsymbol{\beta}}^{\mathrm{T}}\boldsymbol{Z}_j)\widetilde{\boldsymbol{\beta}}^{\mathrm{T}}\boldsymbol{Z}_j\}]$. Define the gradient vector $\nabla l(\boldsymbol{\beta}) = -\partial l_n(\boldsymbol{\beta})/\partial\boldsymbol{\beta}$, the Hessian matrix $\nabla^2 l(\boldsymbol{\beta}) = -\partial^2 l_n(\boldsymbol{\beta})/\partial\boldsymbol{\beta}\partial\boldsymbol{\beta}^{\mathrm{T}}$, the Cholesky decomposition of $\nabla^2 l(\boldsymbol{\beta})$, i.e. $\nabla^2 l(\boldsymbol{\beta}) = \boldsymbol{U}^{\mathrm{T}}\boldsymbol{U}$ and the pseudo response vector $\boldsymbol{R} = (\boldsymbol{U}^{\mathrm{T}})^{-1}\{\nabla^2 l(\boldsymbol{\beta})\boldsymbol{\beta} - \nabla l(\boldsymbol{\beta})\}$. With the second-order Taylor expansion, $-l_n(\boldsymbol{\beta})$ is approximated by $1/2(\boldsymbol{R} - \boldsymbol{U}\boldsymbol{\beta})^{\mathrm{T}}(\boldsymbol{R} - \boldsymbol{U}\boldsymbol{\beta})$. That is, we aim the minimize

$$
1/2(\boldsymbol{R} - \boldsymbol{U}\boldsymbol{\beta})^{\mathrm{T}}(\boldsymbol{R} - \boldsymbol{U}\boldsymbol{\beta}) + n\rho \sum_{j=1}^{q} |\widetilde{\beta}_j|^{-1}|\beta_j|. \tag{5.5}
$$

For any fixed ρ, the following is a complete algorithm for solving (5.4).

Step 1. Obtain $\widetilde{\boldsymbol{\beta}}$, $\widetilde{\boldsymbol{\alpha}}$, $\widetilde{\Lambda}$ and $\widetilde{\psi}$ as the initial estimates using the methods in Wang et. al. (2019), and then compute $\widetilde{R}_i$ $(i = 1, \cdots, n)$.

Step 2. Compute ∇l, $\nabla^2 l$, $\boldsymbol{U}$ and $\boldsymbol{R}$ based on the value of $\widetilde{\boldsymbol{\beta}}$.

Step 3. Minimize the objective function (5.5) with constraint $||\boldsymbol{\beta}|| = 1$.

We propose the criteria for the optimal tuning parameter following steps similar to Fan and Li (2001), Zhang and Lu (2007) and Liu and Zeng (2013). Define the penalty function $p_\rho(|\beta_j|) = \rho|\widetilde{\beta}_j|^{-1}|\beta_j|$. For nonzero β_{j0} and β_j , the derivative of the penalty with respect to β_j can be written as

$$
[p_\rho(|\beta_j|)]' = p'_\rho(|\beta_j|)|\beta_j|^{-1}\beta_j \approx p'_\rho(|\beta_{j0}|)|\beta_{j0}|^{-1}\beta_{j0} = n\rho\beta_{j0}^{-1}.
$$

Therefore, the penalty function can be approximated by

$$p_\rho(|\beta_{j0}|) + 1/2|\beta_{j0}|^{-1}p'_\rho(|\beta_{j0}|)(\beta_j^2 - \beta_{j0}^2),$$

for β_{j0} close to β_j. If $\beta_{j0} = 0$, a consistent initial estimate produces a penalty weight going to infinity and shrinks the adaptive lasso estimate to zero. That is, the objective function can be alternatively approximated by

$$1/2(\boldsymbol{R} - \boldsymbol{U}\boldsymbol{\beta})^{\mathrm{T}}(\boldsymbol{R} - \boldsymbol{U}\boldsymbol{\beta}) + n\rho\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{A}\boldsymbol{\beta},$$

where $\boldsymbol{A} = diag(\widetilde{\beta}_1^{-2}, \cdots, \widetilde{\beta}_q^{-2})$. It is straightforward to obtain the effective number of parameters $e(\rho) = tr(\nabla^2 l(\boldsymbol{\beta}) + \boldsymbol{A})^{-1}\nabla^2 l(\boldsymbol{\beta})$. We can choose the tuning parameter ρ according to the generalized cross validation criteria

$$GCV(\rho) = -ln(\boldsymbol{\beta}_\rho)/[n\{1 - e(\rho)/n\}^2].$$

5.3 Asymptotic Properties

We describe the proposed adaptive Lasso estimator with the penalized objective function based on n samples, which is denoted by $Q_n(\boldsymbol{\beta}) = l_n(\boldsymbol{\beta}) - n\rho_n \sum_{j=1}^q |\beta_j|/|\widetilde{\beta}_j|$. Write the true values of $\boldsymbol{\beta}$ as $\boldsymbol{\beta}_0 = (\boldsymbol{\beta}_{10}, \boldsymbol{\beta}_{20})^{\mathrm{T}}$, where $\boldsymbol{\beta}_{10}$ consists of q_1 zero components and $\boldsymbol{\beta}_{20}$ consists of q_2 nonzero components. Similarly, we define the adaptive lasso estimator $\widehat{\boldsymbol{\beta}}_n = (\widehat{\boldsymbol{\beta}}_{1n}, \widehat{\boldsymbol{\beta}}_{2n})^{\mathrm{T}}$. Also, define the maximum likelihood estimator after variable selection $\check{\boldsymbol{\beta}}_n = (\check{\boldsymbol{\beta}}_{1n}, \check{\boldsymbol{\beta}}_{2n})^{\mathrm{T}}$. Finally, we write the potential interaction covariates $\boldsymbol{Z}$ as $(\boldsymbol{Z}_1^{\mathrm{T}}, \boldsymbol{Z}_2^{\mathrm{T}})$, where $\boldsymbol{Z}_1$ denotes the unimportant covariates and $\boldsymbol{Z}_2$ denotes the important ones. Let Λ_0 denote the true value for the cumulative hazard function. We require the following regularity conditions.

We impose the following regularity conditions.

Condition 1. The true values $\boldsymbol{\alpha}_0$ and $\boldsymbol{\beta}_0$ lie in the interior of a known compact set; $\psi_0(u)$ is a strictly increasing function of u and three-times differentiable in $\mathcal{Z}$, and $\lambda_0(t)$ is twice differentiable in $[0, \tau]$.

Condition 2. The conditional density of C given $\boldsymbol{X}$ has bounded twice differentiable derivative on its support $[0, \tau]$. The conditional distribution of $(A, \boldsymbol{X})$ given $\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z}$ has a continuously differentiable density with respect to a dominating measure.

Condition 3. With probability 1, $P(Y(\tau) = 1|\boldsymbol{X}) > c_0$ for some positive constant c_0. In addition, if $\boldsymbol{c}_1^{\mathrm{T}}\boldsymbol{X} = d_1$ with probability 1 for some constant vector $\boldsymbol{c}_1$ and constant d_1, then $\boldsymbol{c}_1 = \boldsymbol{0}$ and $d_1 = 0$.

Condition 4. The number of the knots satisfies that $K_n \to \infty$ and $n^{-1/2}K_n^7 \to \infty$. In addition, the adjacent distance of the interior knots is between $c^{-1}K_n^{-1}$ and cK_n^{-1} for some positive constant c.

Condition 5. For any positive c_0, $\limsup_{x\to\infty}\{G(c_0x)\}^{-1}\log\{x \sup_{y\leqslant x}G'(y)\} \leqslant 1$.

Remark 1. These are the same set of conditions used in the maximum likelihood estimation in the single-index transformation model. No additional conditions are needed for the adaptive lasso estimator.

Remark 2. The second part of Condition 1 and Condition 2 ensure smoothness for the functions ψ_0 and λ_0, the conditional distribution of C given $\boldsymbol{X}$, and the conditional distribution of $(A, \boldsymbol{X})$ given $\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z}$. Condition 3 ensures that a non-trivial proportion of subjects is censored at τ and that $(1, \boldsymbol{X}^{\mathrm{T}})^{\mathrm{T}}$ is linearly independent. Condition 4 is satisfied by most transformations. For example, the limit goes to zero for the Box-Cox transformation $G(x) = \{(1+x)^\rho - 1\}/\rho$ with $\rho > 0$ and goes to one for the logarithmic transformation $G(x) = r^{-1}\log(1+rx)$.

We state that the adaptive lasso estimator is consistent for the true value at the rate $n^{1/2}$.

Theorem 5.3.1. *Under Conditions 1–5, suppose $n^{1/2}\rho_n = O_p(1)$, then the adaptive lasso estimator satisfies $||\widehat{\boldsymbol{\beta}}_n - \boldsymbol{\beta}_0|| = O_p(n^{-1/2})$.*

The next theorem shows the oracle property of the adaptive lasso estimator. In other words, as sample size increases, the selector keeps the important variables and discard the unimportant ones as if the distinction between the two were known. In addition, the

asymptotic distribution of the important interaction variables is a multivariate normal distribution with the semiparametric efficiency bound attained. The efficiency is because without the penalty, the quadratic approximation to the objective function is exactly the same as the approximate objective function for the single-index transformation model in the M step of an EM algorithm. The penalty is not dominating as shown by the oracle property. Therefore, the estimator attains the same efficiency as in the unpenalized single-index transformation models.

Theorem 5.3.2. *Under Conditions 1–5, suppose* $n^{1/2}\rho_n \to 0$ *and* $n\rho_n \to \infty$*, then under Theorem 1, the adaptive lasso estimator* $\widehat{\boldsymbol{\beta}}_n$ *has the following properties: (i)* $\widehat{\boldsymbol{\beta}}_{1n} = 0$ *with probability tending to 1; (ii)* $n^{1/2}(\widehat{\boldsymbol{\beta}}_{2n} - \boldsymbol{\beta}_{20})$ *converges in distribution to a zero-mean normal random vector with covariance matrix* $\boldsymbol{\Sigma}^{-1}$*, and* $\boldsymbol{\Sigma}^{-1}$ *achieves the semiparametric efficiency bound.*

Last, we show the theoretical properties for the maximum likelihood estimator of the selected important variables when refitting the model without the adaptive lasso penalty after variable selection.

Theorem 5.3.3. *Under Conditions 1–5, the maximum likelihood estimator after variable selection,* $n^{1/2}(\check{\boldsymbol{\beta}}_{2n} - \boldsymbol{\beta}_{20})$ *converges in distribution to a zero-mean normal random vector with covariance matrix* $\boldsymbol{\Sigma}^{-1}$*, and* $\boldsymbol{\Sigma}^{-1}$ *achieves the semiparametric efficiency bound.*

Proofs of the theorems are given in the Appendix.

5.4 Simulation Studies

The simulation study is set up as follows. We considered sample sizes 1000, 2000 and 5000. The baseline hazard function $\Lambda(t)$ was assumed to follow a Weibull distribution with shape parameter 2.5 and scale parameter 2. We considered three link functions: (a) (exponential) $\psi_0(u) = e^u - 0.5$; (b) (linear) $\psi_0(u) = u + 0.4$; and (c) (sine)

$$\psi_0(u) = \begin{cases} -3 \text{ if } u \leqslant -1/2 \\ 2\sin(\pi u) - 1 \text{ if } -1/2 < u \leqslant 1/2 \\ 1 \text{ if } u > 1/2. \end{cases}$$

Table 5.1: Variable selection proportions, the average number of correct and incorrect zero coefficients with the adaptive lasso method.

			Selection Percentage											
r	**Link**	**n**	$\boldsymbol{\beta_{11}}$	$\boldsymbol{\beta_{12}}$	$\boldsymbol{\beta_{13}}$	$\boldsymbol{\beta_{14}}$	$\boldsymbol{\beta_{15}}$	$\boldsymbol{\beta_{16}}$	$\boldsymbol{\beta_{17}}$	$\boldsymbol{\beta_{18}}$	$\boldsymbol{\beta_{19}}$	$\boldsymbol{\beta_{20}}$	**Cor**	**Incr**
0	Expon-	1000	8	6	6	6	7	6	91	93	92	92	14.9	0.3
	ential	2000	2	2	1	2	2	1	100	100	100	100	15.8	0.0
		5000	1	1	1	1	1	1	100	100	100	100	15.9	0.0
	Linear	1000	19	18	15	17	18	20	75	78	78	77	13.0	0.9
		2000	14	14	13	14	13	12	95	97	96	95	13.9	0.2
		5000	5	5	5	6	5	6	100	100	100	100	15.2	0.0
	Sine	1000	3	2	3	3	2	2	92	92	92	92	15.6	0.3
		2000	0	0	0	0	0	0	100	100	100	100	16.0	0.0
		5000	0	0	0	0	0	0	100	100	100	100	16.0	0.0
1	Expon-	1000	11	8	9	9	9	10	73	75	74	74	14.5	1.0
	ential	2000	4	3	3	4	2	2	95	95	96	96	15.6	0.2
		5000	1	2	1	1	1	1	100	100	100	100	15.8	0.0
	Linear	1000	15	16	14	15	16	16	50	53	52	53	13.4	1.9
		2000	13	11	12	13	13	11	77	76	76	75	14.1	1.0
		5000	4	6	5	5	5	5	99	98	98	98	15.2	0.1
	Sine	1000	16	15	15	15	16	15	78	79	78	78	13.6	0.9
		2000	1	1	1	1	1	1	97	97	97	97	15.9	0.1
		5000	0	0	0	0	0	0	100	100	100	100	16.0	0.0

r denotes the transformation parameter in the logarithmic transformation. "Selection percentage" denotes the proportion of replicates where the parameter is selected by the proposed method. "Corr" and "Incr" denotes the average number of correct zero parameters and incorrect zero parameters in the model. Each entry is based on 1000 replicates.

The censoring time was assumed to follow Unif$[0, \tau]$, where τ was chosen such that the censoring rate was 50%. We set the corresponding main effect parameters $(\alpha_{17}, \alpha_{18}, \alpha_{19}, \alpha_{20})^{\mathrm{T}}$ to $(-0.4, -0.2, 0.2, 0.4)^{\mathrm{T}}$ and $\alpha_j = 0$ for $j = 1 \cdots 16$. Similarly, we set the interaction effect parameters $(\beta_{17}, \beta_{18}, \beta_{19}, \beta_{20})^{\mathrm{T}}$ to $(-0.5, -0.5, 0.5, 0.5)^{\mathrm{T}}$, and $\beta_j = 0$ for $j = 1 \cdots 16$. The treatment assignment A was assumed to be independent of $\boldsymbol{X}, \boldsymbol{Z}$ and to follow a Bernoulli

distribution with success probability of 0.5. We simulated 1000 replicates.

Table 5.2: Simulation results on $\beta_{17} - \beta_{20}$ with exponential link.

			Before selection		**After selection**			
r	**n**		**Bias**	**SE**	**Bias**	**SE**	**SEE**	**CP**
0	1000	β_{17}	0.07	0.11	0.01	0.11	0.09	0.87
		β_{18}	0.07	0.11	0.01	0.11	0.09	0.87
		β_{19}	-0.07	0.11	-0.01	0.11	0.09	0.87
		β_{20}	-0.06	0.11	-0.01	0.11	0.09	0.89
	2000	β_{17}	0.03	0.07	0.00	0.07	0.06	0.93
		β_{18}	0.02	0.07	0.00	0.07	0.06	0.94
		β_{19}	-0.02	0.07	0.00	0.07	0.06	0.94
		β_{20}	-0.02	0.07	0.00	0.06	0.07	0.95
	5000	β_{17}	0.01	0.04	0.00	0.04	0.04	0.95
		β_{18}	0.01	0.04	0.00	0.04	0.04	0.94
		β_{19}	-0.01	0.04	0.00	0.04	0.04	0.94
		β_{20}	-0.01	0.04	0.00	0.04	0.04	0.94
1	1000	β_{17}	0.12	0.15	0.04	0.15	0.11	0.81
		β_{18}	0.12	0.14	0.05	0.14	0.11	0.86
		β_{19}	-0.12	0.15	-0.04	0.15	0.12	0.83
		β_{20}	-0.11	0.15	-0.03	0.16	0.12	0.86
	2000	β_{17}	0.05	0.10	-0.02	0.09	0.09	0.95
		β_{18}	0.05	0.10	-0.01	0.10	0.09	0.92
		β_{19}	-0.05	0.10	0.01	0.10	0.09	0.92
		β_{20}	-0.04	0.10	0.01	0.10	0.10	0.92
	5000	β_{17}	0.01	0.06	-0.01	0.06	0.06	0.96
		β_{18}	0.02	0.06	0.01	0.06	0.06	0.94
		β_{19}	-0.03	0.06	-0.01	0.06	0.06	0.95
		β_{20}	-0.01	0.06	0.01	0.06	0.06	0.95

r denotes the transformation parameter in the logarithmic transformation. Bias and SE are the bias and standard error of the parameter estimator, respectively; SEE is the mean of the standard error estimator; CP is the coverage probability of the 95% confidence interval. Each entry is based on 1000 replicates.

We obtained the initial estimates using the expectation-maximization algorithm described in section 5.2.2. We then implemented the adaptive lasso variable selection procedure in section 5.2.3 to select important features in the optimal treatment rule estimation. We considered the grid points $(2^{-8}, 2^{-7}, \ldots, 1, \ldots, 2^7, 2^8)$ for tuning parameter ρ, and reported

the results with the optimal value of ρ. Table 5.1 summarizes the proportion of times each parameter is selected in the adaptive lasso estimation based on 1000 simulated replicates for $\beta_{11} - \beta_{20}$. With each transformation function and link function, the proposed method is able to select the correct set of parameters when sample size increases. Transformation model with $r = 1$ requires a larger sample size for the method to correctly select the important variables compared to $r = 0$.

Table 5.3: Inference results on $\beta_{17} - \beta_{20}$ with linear link.

			Before selection		**After selection**			
r	**n**		**Bias**	**SE**	**Bias**	**SE**	**SEE**	**CP**
0	1000	β_{17}	0.13	0.15	0.05	0.16	0.11	0.78
		β_{18}	0.12	0.14	0.06	0.15	0.11	0.82
		β_{19}	-0.13	0.15	-0.06	0.15	0.11	0.81
		β_{20}	-0.13	0.14	-0.06	0.14	0.12	0.85
	2000	β_{17}	0.06	0.10	0.00	0.10	0.09	0.92
		β_{18}	0.07	0.10	0.00	0.11	0.09	0.88
		β_{19}	-0.06	0.10	0.00	0.11	0.09	0.91
		β_{20}	-0.07	0.10	-0.01	0.10	0.09	0.92
	5000	β_{17}	0.02	0.06	0.00	0.06	0.06	0.95
		β_{18}	0.02	0.06	0.00	0.06	0.06	0.95
		β_{19}	-0.02	0.06	0.00	0.06	0.06	0.94
		β_{20}	-0.02	0.06	0.00	0.06	0.06	0.95
1	1000	β_{17}	0.19	0.17	0.10	0.17	0.14	0.80
		β_{18}	0.19	0.17	0.12	0.18	0.14	0.78
		β_{19}	-0.20	0.17	-0.12	0.18	0.14	0.80
		β_{20}	-0.19	0.17	-0.11	0.17	0.23	0.86
	2000	β_{17}	0.11	0.14	0.03	0.15	0.12	0.87
		β_{18}	0.12	0.14	0.04	0.15	0.12	0.86
		β_{19}	-0.12	0.14	-0.03	0.14	0.12	0.86
		β_{20}	-0.11	0.14	-0.03	0.14	0.12	0.91
	5000	β_{17}	0.04	0.08	-0.01	0.09	0.09	0.94
		β_{18}	0.05	0.09	0.00	0.08	0.09	0.96
		β_{19}	-0.05	0.09	0.00	0.09	0.09	0.95
		β_{20}	-0.04	0.09	0.01	0.09	0.09	0.95

See note to table 5.2.

After the variable selection procedure, we obtained the maximum likelihood estimator by refitting the single-index transformation model using the expectation-maximization algorithm

Table 5.4: Inference results on $\beta_{17} - \beta_{20}$ with sine link.

			Before selection		After selection			
r	**n**		**Bias**	**SE**	**Bias**	**SE**	**SEE**	**CP**
0	1000	β_{17}	0.01	0.04	0.00	0.03	0.03	0.95
		β_{18}	0.01	0.04	0.00	0.03	0.03	0.96
		β_{19}	-0.01	0.04	0.00	0.03	0.03	0.93
		β_{20}	-0.01	0.04	0.00	0.03	0.04	0.96
	2000	β_{17}	0.00	0.02	0.00	0.02	0.02	0.95
		β_{18}	0.00	0.02	0.00	0.02	0.02	0.95
		β_{19}	0.00	0.02	0.00	0.02	0.02	0.95
		β_{20}	-0.01	0.02	0.00	0.02	0.02	0.96
	5000	β_{17}	0.00	0.01	0.00	0.01	0.01	0.94
		β_{18}	0.00	0.01	0.00	0.01	0.01	0.96
		β_{19}	0.00	0.01	0.00	0.01	0.01	0.94
		β_{20}	0.00	0.01	0.00	0.01	0.01	0.95
1	1000	β_{17}	0.03	0.06	0.02	0.06	0.05	0.93
		β_{18}	0.03	0.06	0.01	0.06	0.05	0.94
		β_{19}	-0.03	0.07	-0.01	0.06	0.05	0.92
		β_{20}	-0.03	0.06	-0.01	0.06	0.06	0.93
	2000	β_{17}	0.01	0.04	0.00	0.03	0.04	0.97
		β_{18}	0.01	0.04	0.00	0.03	0.04	0.97
		β_{19}	-0.01	0.04	0.00	0.03	0.04	0.97
		β_{20}	-0.01	0.04	0.00	0.03	0.04	0.98
	5000	β_{17}	0.00	0.02	0.00	0.02	0.02	0.98
		β_{18}	0.00	0.02	0.00	0.02	0.02	0.97
		β_{19}	0.00	0.02	0.00	0.02	0.02	0.97
		β_{20}	0.00	0.02	0.00	0.02	0.02	0.98

See note to table 5.2.

proposed in section 5.2.2 on the set of parameters that are selected using the adaptive Lasso method. The results concerning the updated estimates and the corresponding variance estimator are summarized in tables 5.2-5.4. With all link functions, the bias in the original estimate can be reduced considerably by refitting the model, especially for the sample size of 1000. When sample size is 2000 or larger, the biases with the maximum likelihood estimators are small and the standard error estimates are close to the sample standard errors. The 95% confidence intervals based on the estimated parameters and standard error estimates have accurate coverage for true parameters.

5.5 Application to evaluate adjuvant therapy of large-bowel carcinoma

We applied the proposed method to the large-bowel carcinoma clinical trial (Moertel et al., 1990; Laurie et al., 1989). This is the same dataset as analyzed in section 4.5. The aim of the clinical trial was to assess the effectiveness of fluorouracil (5-FU) plus levamisole versus levamisole alone and the observation group.

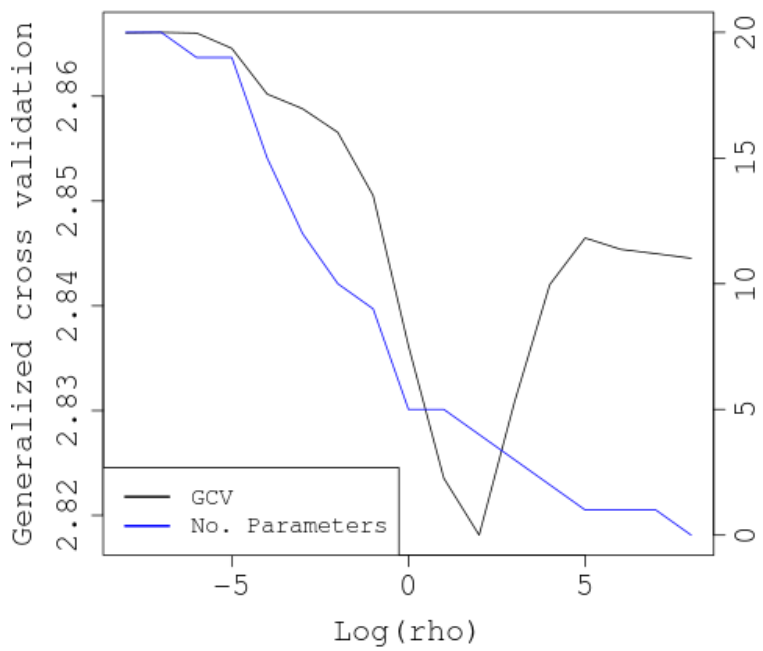

Figure 5.1: Details of the generalized cross validation with the adaptive lasso approach for variable selection.

There are a total of 20 variables measured in the large-bowel carcinoma data: gender, obstruction by the lesion, development of cancer cells (grade 1; grade 2-3; grade 4; others), location of primary neoplasm (cecum; sigmoid colon; others), extent of local speed (submucosa; muscular; serosa; others), histologic type (adenocarcinoma; colloid; signet ring; others), perforation, adherence, regional implants, number of positive nodes, age, number of days from tumor resection to start of treatment and pre-operative CEA level. In the exploratory analysis described in section 4.5, we fitted the saturated proportional hazards model including

linear interaction between treatment and covariates. A backward selection procedure was implemented to determine the variables to be included in the model. In this analysis, we applied the adaptive lasso method to the same dataset for the variable selection purpose.

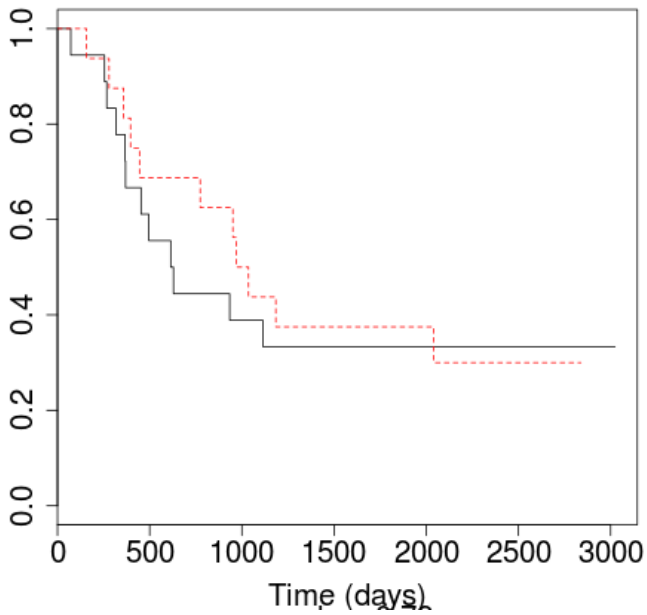

Figure 5.2: Estimation of the probability of disease relapse for patients who receive a different treatment recommendation under the two models. The black and red curves pertain, respectively, to patients who follow the treatments assigned with backward selection variables and adaptive lasso variables.

We fitted the transformation single index model with the logarithmic transformation parameter $r = 1.5$ to obtain the initial estimate. Then we used the algorithm described in section 5.2.3 to select important variables with the tuning parameter taking values from the grid points $(2^{-8}, 2^{-7}, \ldots, 1, \ldots, 2^{7}, 2^{8})$. According to the generalized cross validation criteria, $\rho = 4$ is chosen as the optimal tuning parameter. Details of the generalized cross validation criteria can be found in Figure 5.1. Variables representing the extent of local speed (submucosa; muscular; serosa; others) and age are selected in this model. We applied the transformation single-index model described in section 5.2.2 to the selected features and

obtained the post-selection maximum likelihood estimator and the optimal treatment rule.

Table 5.5: Estimation of the parameters under different transformation models

Before selection		After selection			Original		
Parameters	**Est.**	**Est.**	**Std. Err.**	***P*-value**	**Est.**	**Std. Err.**	***P*-value**
Interactions							
Gender	−0.02						
Obstruction	0.00						
Grade 1 development	0.00						
Grade 2-3 development	0.01						
Grade 4 development	−0.27						
Cecum location	−0.01						
Sigmoid location	0.01						
Submucosa speed	0.07	−0.64	0.29	0.01			
Muscular speed	0.07	−0.58	0.27	0.01			
Serosa speed	0.08	−0.47	0.47	0.16			
Adenocarcinoma type	0.00						
Colloid type	−0.56						
Signet ring type	−0.67						
Perforation	−0.01				0.04	1.63	0.98
Adherence	0.05				0.65	0.36	0.07
Regional implants	−0.01						
Log positive nodes	0.04				0.24	0.23	0.30
Days before treatment	0.00						
Pre-operative CEA level	0.00				0.28	0.25	0.26
Age	−0.10	−0.17	0.50	0.37	-0.67	0.30	0.03

"After selection" denotes the estimation results by refitting the transformation single-index model on variables selected by the adaptive lasso method, $\rho = 4$. "Original" denotes the estimation results as described in section 4.5, where the features were selected using a backward selection procedure. "Est." and "Std. Err." denote the estimate and the standard error estimate for a parameter.

In Table 5.5, we reported the estimation results for the treatment-covariate interaction effect with the adaptive lasso approach where $\rho = 1$. Results with the backward selection approach was included as a comparison. The variables selected by the adaptive lasso method is quite different from those in the backward selection results. Age is a common feature that is selected under both models. The three variables representing the extent of local speed are deemed as important by the adaptive lasso penalty. On the other hand, backaward selection result favors perforation, adherence, number of positive nodes and the pre-operative CEA

level.

We compared the treatment rule estimated with different variable selection approaches. Table 5.6 shows that, out of 908 patients, 63 would receive different treatments using the adaptive lasso variable selection method compared to the original results using the backward selection approach. Figure 5.2 shows that probabilities of disease relapse decrease if patients follow the treatment regimen recommended by refitting the transformation single-index model on variables selected with the adaptive lasso procedure.

Table 5.6: Treatment assignments with the adaptive lasso procedure and with the original backward selection approach

	Original		
Adaptive Lasso	**5-FU + Levamisole**	**Observation/ Levamisole**	**Total**
5-FU + Levamisole	834	37	871
Observation/Levamisole	26	11	37
Total	860	48	908

5.6 Discussion

In this chapter, we are mainly concerned with the variable selection in the treatment-covariate interaction term. The method could be easily extended to select the important main effect and interaction effect of covariates at the same time. The penalized estimation procedure can be generalized to other commonly used penalties. The asymptotic properties, such as the oracle property, the asymptotic normal distribution and the efficiency of the interaction term remain valid with both extensions.

With the advance of technology in bioinformatics, genetic information on patients, such as genotype and protein expression data, can be obtained with a reasonable cost. Such genetic dataset usually contains hundreds of thousands variables for each patient. An important extension of this chapter is to develop a variable selection method for the single-index transformation model with $p >> n$. The initial estimate could be obtained by, for example, a penalized estimator with the ridge penalty. The estimation procedure and the asymptotic

properties can be developed in a similar fashion with this new initial estimate

CHAPTER 6: SUMMARY AND FUTURE RESEARCH

In this dissertation, we have studied single-index models for estimation of the optimal treatment regime with censored outcomes. The proposed model allows for flexible non-linear treatment-covariate interactions and yet the treatment rule remains a simple linear form. This research enables individualized treatment options to be assigned to patients by incorporating their features into the decision process so that each patient receives the maximum benefit on the individual level.

In Chapter 3, we proposed to apply the single-index function to the Cox proportional hazards model in order to account for the treatment covariate interaction in a flexible manner. We designed a non-parametric maximum likelihood estimation procedure and showed the theoretical properties for the parametric and nonparametric estimators. We showed through extensive comparisons that the proposed method outperform the value-based optimal treatment estimation approach both in simulation studies and in a data application to the ACTG data.

In Chapter 4, we relaxed the proportional hazards assumption and extended the single-index framework to transformation models. The transformation function presents new challenges in both the asymptotic properties and the estimation procedures. We designed an expectation-maximization algorithm that greatly reduces the computational burden. Asymptotic distribution and the semiparametric efficiency of parametric component in the model were shown by applying the theory of empirical process and semiparametric theory. We demonstrated the performance of the proposed method in simulations with different link and transformation functions. In a data application, we showed the proposed method delayed disease relapse for patients in a large-bowel carcinoma clinical trial.

In Chapter 5, we investigated the variable selection procedure in applying the single-index

model for estimation of the optimal treatment rule. We proposed an adaptive lasso approach, showed its oracle property as well as the asymptotic behaviour of the maximum likelihood estimator after selection. Finite sample performance was studied by simulation studies under various settings. The large-bowel carcinoma data was analyzed to illustrate the performance of the treatment assignment resulting from the adaptive lasso variable selection.

Several future directions of research naturally extend from this dissertation. One question that merits future investigation is the relationship between value inference and the theoretical properties of the parameters in the single-index model. Researchers may be interested in quantities that summarize patient performance, such as the survival function under the treatment rule, or the proportion of patients who receive the correct optimal treatment assignment. These quantities can be estimated by the plug-in estimator, and their asymptotic properties follow from the delta method and the limiting distributions of $\widehat{\boldsymbol{\alpha}}$, $\widehat{\boldsymbol{\beta}}$, and $\widehat{\psi}$. We have already proved that the parametric component, namely $\widehat{\boldsymbol{\alpha}}$ and $\widehat{\boldsymbol{\beta}}$, are semiparametrically efficient. Based on these findings, it would be interesting to see if we could establish the efficiency for the value function.

Another enticing possibility is to extend the proposed method to a more complex setting, such as multiple decision points and multiple treatment options. Compared to the many works on personalized medicine, there are few literature available on the optimal treatment estimation with more than one decision point. Jiang et al. (2017) proposed to model the censoring distribution and used the inverse probability weighting method to estimate the patient survival probability if he followed multiple treatment assignment. For multiple treatment options, we can include multiple index functions in the model, with each index function representing the treatment rule for one option.

Last but not least, it is of great importance to develop model diagnosis criteria. Some robust analysis have been done to demonstrate the robustness of the proposed model. For example, in Chapter 3, we investigated the behavior of the single-index proportional hazards model when there are in fact two indices in the model but we only fit with on index function.

We also considered the results when the transformation parameter is misspecified. In both scenarios, we showed that the estimation results remained robust in terms of the treatment recommendation. Nevertheless, it is essential to find out what model assumptions are the treatment assignment sensitive to, and design statistics measuring the departure from the proposed model.

APPENDIX : PROOFS

A.1 Proof for Chapter 3

A.1.1 Proof of Theorem 3.3.1

According to Theorem 6.25 of Schumaker (1981), there exists a function $\widetilde{\psi}(u) = \sum_{k=1}^{K_n} \widetilde{\gamma}_k B_k(u)$ such that $||\widetilde{\psi} - \psi_0||_{W^{1,\infty}(\mathcal{Z})} \leqslant O(K_n^{-2})$ and $||\widetilde{\psi} - \psi_0||_{L^2(\mathcal{Z})} \leqslant O(K_n^{-7/2})$. Since ψ_0 is strictly increasing, it follows that $\widetilde{\psi}'(u) > 0$ for all $u \in \mathcal{Z}$ when n large enough.

For any $\boldsymbol{\alpha}$ and η, define

$$\vartheta_n(\boldsymbol{\alpha}, \eta) = \Delta\Bigg(\boldsymbol{\alpha}^{\mathrm{T}}\boldsymbol{X} + A\eta(\boldsymbol{Z}) - \log\Bigg[n^{-1}\sum_j \exp\{\boldsymbol{\alpha}^{\mathrm{T}}\boldsymbol{X}_j + A_j\eta(\boldsymbol{Z}_j)\}Y_j(\widetilde{T})\Bigg]\Bigg),$$

and

$$\vartheta_n^*(\boldsymbol{\alpha}, \eta) = \Delta\big(\boldsymbol{\alpha}^{\mathrm{T}}\boldsymbol{X} + A\eta(\boldsymbol{Z}) - \log\big[E\{\exp\{\boldsymbol{\alpha}^{\mathrm{T}}\boldsymbol{X} + A\eta(\boldsymbol{Z})\}Y(\widetilde{T})\}\big]\big).$$

Also, let $\eta_0(\boldsymbol{Z}) = \psi_0(\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z})$, $\widehat{\eta}(\boldsymbol{Z}) = \widehat{\psi}(\widehat{\boldsymbol{\beta}}^{\mathrm{T}}\boldsymbol{Z})$, and $\widetilde{\eta}(\boldsymbol{Z}) = \widetilde{\psi}(\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z})$. Because the profile log-likelihood (3.2) is maximized at $\widehat{\boldsymbol{\alpha}}, \widehat{\boldsymbol{\beta}}$, and $\widehat{\psi}$, we have

$$\mathbb{P}_n\vartheta_n(\widehat{\boldsymbol{\alpha}}, \widehat{\eta}) \geqslant \mathbb{P}_n\vartheta_n(\boldsymbol{\alpha}_0, \widetilde{\eta}). \tag{1}$$

Consider the class of functions

$$\mathcal{F} = \Bigg\{\exp\{\boldsymbol{\alpha}^{\mathrm{T}}\boldsymbol{X} + A\eta(\boldsymbol{Z})\}Y(t) : \boldsymbol{\alpha} \in \Theta_1, ||\boldsymbol{\beta}|| = 1, \\ \eta(\boldsymbol{u}) = \sum_{k=1}^{K_n+1} \gamma_k B_k(\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{u}) \text{ with} - M_n \leqslant \gamma_1 \leqslant \cdots \leqslant \gamma_{K_n+1} \leqslant M_n\Bigg\},$$

where Θ_1 is a compact set containing $\boldsymbol{\alpha}_0$ as given in Condition 3.1. This class has an envelope function of order $\exp(M_n)$. In addition, both $\{\boldsymbol{\alpha}^{\mathrm{T}}\boldsymbol{X}\}$ and $\{\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z}\}$ belong to the bounded finite-dimensional space and thus are Vapnik-Chervonenkis (VC)-major classes (Vapnik and

Chervonenkis, 1971). Since $\sum_{k=1}^{K_n+1} \gamma_k B_k(u)$ is monotone and bounded by $O(M_n)$,

$$\left\{ M_n^{-1}\eta(\boldsymbol{Z}) : \eta(\boldsymbol{Z}) = \sum_{k=1}^{K_n+1} \gamma_k B_k(\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z}) \text{ with} - M_n \leqslant \gamma_1 \leqslant \cdots \leqslant \gamma_{K_n+1} \leqslant M_n \right\}$$

is a bounded VC-major class (van der Vaart and Wellner, 1996, Lemma 2.6.19). Since any function in $\mathcal{F}$ is Lipschitz-continuous with respect to $\boldsymbol{\alpha}^{\mathrm{T}}\boldsymbol{X}$ and $\eta(\boldsymbol{Z})$, we see that, for any probability measure Q, the ϵ-covering number for $\mathcal{F}$ satisfies that

$$\log N(\epsilon, \mathcal{F}, L_2(Q)) \leqslant c_1(\epsilon^{-1}\exp(M_n)M_n)^{2v_1/(v_1+2)}$$

for some positive constants c_1 and v_1 (van der Vaart and Wellner, 1996, Theorem 2.6.9). By Theorem 2.14.1 of van der Vaart and Wellner (1996), uniformly in $t \in [0, \tau]$,

$$(\mathbb{P}_n - \mathbb{P})\big[\exp\big\{\widehat{\boldsymbol{\alpha}}^{\mathrm{T}}\boldsymbol{X} + A\widehat{\eta}(\boldsymbol{Z})\big\}Y(t)\big] = O_P(M_n \exp(M_n)n^{-1/2}).$$

Clearly,

$$(\mathbb{P}_n - \mathbb{P})\left[\exp\left\{\boldsymbol{\alpha}_0^{\mathrm{T}}\boldsymbol{X} + A\widetilde{\eta}(\boldsymbol{Z})\right\}Y(t)\right] = O_P(n^{-1/2}).$$

Under Conditions 3.1 and 3.3, $\mathbb{P}\big[\exp\big\{\widehat{\boldsymbol{\alpha}}^{\mathrm{T}}\boldsymbol{X} + A\widehat{\eta}(\boldsymbol{Z})\big\}Y(t)\big]$ is bounded by $O(\exp(-M_n))$ from below for all $t \in [0, \tau]$. Thus, uniformly in $(\Delta, \boldsymbol{X}, A, \tilde{T})$,

$$\vartheta_n(\widehat{\boldsymbol{\alpha}}, \widehat{\eta}) - \vartheta_n^*(\widehat{\boldsymbol{\alpha}}, \widehat{\eta}) = O_p(M_n \exp(2M_n)n^{-1/2}), \tag{2}$$

and

$$\vartheta_n(\boldsymbol{\alpha}_0, \widetilde{\eta}) - \vartheta_n^*(\boldsymbol{\alpha}_0, \widetilde{\eta}) = O_p(M_n \exp(2M_n)n^{-1/2}). \tag{3}$$

Combining (1), (2) and (3), we obtain

$$(\mathbb{P}_n - \mathbb{P})\left\{\vartheta_n^*(\widehat{\boldsymbol{\alpha}}, \widehat{\eta}) - \vartheta_n^*(\boldsymbol{\alpha}_0, \widetilde{\eta})\right\} + O_p(M_n \exp(2M_n) n^{-1/2}) \geqslant$$
$$- \mathbb{P}\{\vartheta_n^*(\widehat{\boldsymbol{\alpha}}, \widehat{\eta}) - \vartheta_n^*(\boldsymbol{\alpha}_0, \widetilde{\eta})\}. \quad (4)$$

By similar arguments for showing (2), the first term on the left side of (4) is bounded by $O_P(M_n \exp(2M_n) n^{-1/2})$. On the right side of (4), we take the second-order Taylor expansion of $\vartheta_n^*(\widehat{\boldsymbol{\alpha}}, \widehat{\eta})$ around $(\boldsymbol{\alpha}_0, \eta_0)$ and note that the Hessian matrix of the partial log-likelihood function is strictly negative with its eigenvalues bounded by $-c_0 \exp(-2M_n)$ for some constant $c_0 > 0$. Thus

$$\begin{aligned}
&\mathbb{P}\vartheta_n^*(\widehat{\boldsymbol{\alpha}}, \widehat{\eta}) - \mathbb{P}\vartheta_n^*(\boldsymbol{\alpha}_0, \widetilde{\eta}) \\
=&\mathbb{P}\vartheta_n^*(\widehat{\boldsymbol{\alpha}}, \widehat{\eta}) - \mathbb{P}\vartheta_n^*(\boldsymbol{\alpha}_0, \eta_0) + O_p(K_n^{-7}) \\
\leqslant& - c_0 \exp(-2M_n) E\left[||\widehat{\boldsymbol{\alpha}} - \boldsymbol{\alpha}_0||^2 + \{\widehat{\eta}(\boldsymbol{Z}) - \widetilde{\eta}(\boldsymbol{Z}) + \widetilde{\eta}(\boldsymbol{Z}) - \eta_0(\boldsymbol{Z})\}^2\right] + O_p(K_n^{-7}) \\
\leqslant& - c_0 \exp(-2M_n) E\left[||\widehat{\boldsymbol{\alpha}} - \boldsymbol{\alpha}_0||^2 + \{\widehat{\eta}(\boldsymbol{Z}) - \widetilde{\eta}(\boldsymbol{Z})\}^2/2 - \{\widetilde{\eta}(\boldsymbol{Z}) - \eta_0(\boldsymbol{Z})\}^2\right] + O_p(K_n^{-7}).
\end{aligned}$$

As a result,

$$O_p(M_n \exp(4M_n) n^{-1/2}) + O_p(\exp(2M_n) K_n^{-7})$$
$$\geqslant c_0 E\left\{||\widehat{\boldsymbol{\alpha}} - \boldsymbol{\alpha}_0||^2 + |\widehat{\eta}(\boldsymbol{Z}) - \widetilde{\eta}(\boldsymbol{Z})|^2/2\right\}. \quad (5)$$

By the choices of M_n and K_n, we conclude that $\widehat{\boldsymbol{\alpha}} \to_p \boldsymbol{\alpha}_0$, and

$$O_p(M_n \exp(4M_n) n^{-1/2}) + O_p(\exp(2M_n) K_n^{-7}) \geqslant E[|\widehat{\eta}(\boldsymbol{Z}) - \widetilde{\eta}(\boldsymbol{Z})|^2].$$

Recall that $\widehat{\eta}(\boldsymbol{Z}) = \widehat{\psi}(\widehat{\boldsymbol{\beta}}^{\mathrm{T}} \boldsymbol{Z})$ and $\widetilde{\eta}(\boldsymbol{Z}) = \widetilde{\psi}(\boldsymbol{\beta}_0^{\mathrm{T}} \boldsymbol{Z})$. Since $E\left[\left\{\widehat{\psi}(\widehat{\boldsymbol{\beta}}^{\mathrm{T}} \boldsymbol{Z}) - \widetilde{\psi}(\boldsymbol{\beta}_0^{\mathrm{T}} \boldsymbol{Z})\right\}^2\right] = E\left(\left[\widehat{\psi}(\widehat{\boldsymbol{\beta}}^{\mathrm{T}} \boldsymbol{Z}) - E\left\{\widetilde{\psi}(\boldsymbol{\beta}_0^{\mathrm{T}} \boldsymbol{Z}) | \widehat{\boldsymbol{\beta}}^{\mathrm{T}} \boldsymbol{Z}\right\}\right]^2\right) + E\left(\left[E\left\{\widetilde{\psi}(\boldsymbol{\beta}_0^{\mathrm{T}} \boldsymbol{Z}) | \widehat{\boldsymbol{\beta}}^{\mathrm{T}} \boldsymbol{Z}\right\} - \widetilde{\psi}(\boldsymbol{\beta}_0^{\mathrm{T}} \boldsymbol{Z})\right]^2\right)$, inequality

(5) yields

$$O_p(M_n \exp(4M_n)n^{-1/2}) + O_p(\exp(2M_n)K_n^{-7}) \geqslant \\ E\left(\left[E\left\{\psi_0(\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z})|\widehat{\boldsymbol{\beta}}^{\mathrm{T}}\boldsymbol{Z}\right\} - \psi_0(\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z})\right]^2\right).$$

Because $\widehat{\boldsymbol{\beta}}$ is bounded, any subsequence of $\{\widehat{\boldsymbol{\beta}}_n\}$ has a convergent sub-subsequence with a limit denoted by $\boldsymbol{\beta}^*$. By Condition 3.1 and the continuous mapping theorem, $\psi_0^{-1}(E\{\psi_0(\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z})|\boldsymbol{\beta}^{*\mathrm{T}}\boldsymbol{Z}\}) = \boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z}$ almost surely. Condition 3.4 then entails that $\boldsymbol{\beta}^* = \boldsymbol{\beta}_0$. Therefore, $\widehat{\boldsymbol{\beta}} \to_p \boldsymbol{\beta}_0$. It also follows from Condition 3.4 that

$$O_p(M_n \exp(4M_n)n^{-1/2}) + O_p(\exp(2M_n)K_n^{-7}) \geqslant ||\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0||^2. \tag{6}$$

On the other hand,

$$|(\widehat{\psi}(\widehat{\boldsymbol{\beta}}^{\mathrm{T}}\boldsymbol{Z}) - \widetilde{\psi}(\widehat{\boldsymbol{\beta}}^{\mathrm{T}}\boldsymbol{Z}))|^2 \leqslant 2|(\widehat{\psi}(\widehat{\boldsymbol{\beta}}^{\mathrm{T}}\boldsymbol{Z}) - \widetilde{\psi}(\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z}))|^2 + 2|(\widetilde{\psi}(\widehat{\boldsymbol{\beta}}^{\mathrm{T}}\boldsymbol{Z}) - \widetilde{\psi}(\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z})|^2.$$

It follows from (5) that

$$O_p(M_n \exp(4M_n)n^{-1/2}) + O_p(\exp(2M_n)K_n^{-7}) \\ \geqslant E\big\{|(\widehat{\psi}(\widehat{\boldsymbol{\beta}}^{\mathrm{T}}\boldsymbol{Z}) - \widetilde{\psi}(\widehat{\boldsymbol{\beta}}^{\mathrm{T}}\boldsymbol{Z})|^2 - 2|(\widetilde{\psi}(\widehat{\boldsymbol{\beta}}^{\mathrm{T}}\boldsymbol{Z}) - \widetilde{\psi}(\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z})|^2)\big\}.$$

In light of (6),

$$O_p(K_n^2 M_n \exp(4M_n)n^{-1/2}) + O_p(\exp(2M_n)K_n^{-5}) \geqslant E\big\{|\widehat{\psi}(\widehat{\boldsymbol{\beta}}^{\mathrm{T}}\boldsymbol{Z}) - \widetilde{\psi}(\widehat{\boldsymbol{\beta}}^{\mathrm{T}}\boldsymbol{Z})|^2\big\}.$$

Because the L_2-norm between two functions $\widehat{\psi}$ and $\widetilde{\psi}$ is bounded from below by the Euclidean norm of the corresponding coefficient vectors subject to a constant (De Boor, 1978, p. 155),

we obtain

$$O_p(K_n^2 M_n \exp(4M_n) n^{-1/2}) + O_p(\exp(2M_n) K_n^{-5}) \geqslant \sum_{k=1}^{K_n} |\widehat{\gamma}_k - \widetilde{\gamma}_k|^2.$$

It follows that

$$\begin{aligned}
&||\widehat{\psi} - \widetilde{\psi}||_{W^{1,\infty}(\mathcal{Z})} \qquad (7)\\
\leqslant &\sum_{k=1}^{K_n} |\widehat{\gamma}_k - \widetilde{\gamma}_k| \; ||B_k||_\infty + \sum_{k=1}^{K_n} |\widehat{\gamma}_k - \widetilde{\gamma}_k| \; ||B_k'||_\infty \\
\leqslant &K_n^{1/2} \left\{ \sum_{k=1}^{K_n} |\widehat{\gamma}_k - \widetilde{\gamma}_k|^2 \right\}^{1/2} ||B_k||_\infty + K_n^{1/2} \left\{ \sum_{k=1}^{K_n} |\widehat{\gamma}_k - \widetilde{\gamma}_k|^2 \right\}^{1/2} ||B_k'||_\infty \\
= &O_p\left(K_n^2 M_n^{1/2} \exp(2M_n) n^{-1/4} + \exp(M_n) K_n^{-1/2}\right).
\end{aligned}$$

By the choices of K_n and M_n, we obtain $||\widehat{\psi} - \widetilde{\psi}||_{W^{1,\infty}(\mathcal{Z})} \to 0$.

To determine the convergence rate, we repeat the above arguments but notice that the left side of (5) becomes $o_p(n^{-1/2}) + O_p(K_n^{-7})$ because of the convergence of $\widehat{\boldsymbol{\alpha}}$, $\widehat{\boldsymbol{\beta}}$, and $\widehat{\psi}$. By Condition 3.5, we obtain $||\widehat{\boldsymbol{\alpha}} - \boldsymbol{\alpha}_0||^2 + ||\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0||^2 + ||\widehat{\psi} - \psi_0||^2_{L_2(\mathcal{Z})} = o_p(n^{-1/2})$.

A.1.2 Proof of Theorem 3.3.2

Write $\boldsymbol{\theta} = (\boldsymbol{\alpha}^{\mathrm{T}}, \boldsymbol{\beta}_{-q}^{\mathrm{T}})^{\mathrm{T}}$, $\boldsymbol{\theta}_0 = (\boldsymbol{\alpha}_0^{\mathrm{T}}, \boldsymbol{\beta}_{0,-q}^{\mathrm{T}})^{\mathrm{T}}$, and $\boldsymbol{\phi} = (\lambda, \psi)^{\mathrm{T}}$. We wish to obtain the least favorable direction for $\boldsymbol{\theta}$, defined as a function $(\boldsymbol{h}_\lambda^{\mathrm{T}}, \boldsymbol{h}_\psi^{\mathrm{T}})^{\mathrm{T}}$, such that $E[l_{\boldsymbol{\phi}}^* l_{\boldsymbol{\theta}}] = E[l_{\boldsymbol{\phi}}^* l_{\boldsymbol{\phi}}[(\boldsymbol{h}_\lambda^{\mathrm{T}}, \boldsymbol{h}_\psi^{\mathrm{T}})^{\mathrm{T}}]]$, where $l_{\boldsymbol{\theta}}$ is the score for $\boldsymbol{\theta}$, $l_{\boldsymbol{\phi}}[(\boldsymbol{h}_\lambda^{\mathrm{T}}, \boldsymbol{h}_\psi^{\mathrm{T}})^{\mathrm{T}}]$ is the score function for $\boldsymbol{\phi}$ along the submodel $\boldsymbol{\phi} + \epsilon(\boldsymbol{h}_\lambda^{\mathrm{T}}, \boldsymbol{h}_\psi^{\mathrm{T}})^{\mathrm{T}}$, and $l_{\boldsymbol{\phi}}^*$ is the dual operator of $l_{\boldsymbol{\phi}}$. Equivalently, $l_{\boldsymbol{\phi}}[(\boldsymbol{h}_\lambda^{\mathrm{T}}, \boldsymbol{h}_\psi^{\mathrm{T}})^{\mathrm{T}}]$ is the projection of $l_{\boldsymbol{\theta}}$ on the tangent space spanned by the score functions for (λ, ψ).

Let Υ_λ and Υ_ψ denote the nuisance tangent space for λ and ψ, respectively. By direct calculation, the score function of $\boldsymbol{\theta}$ in model (2.1) is given by

$$l_{\boldsymbol{\theta}} = \int \boldsymbol{g}(\boldsymbol{X}, A) dM(u, \boldsymbol{X}, A).$$

According to Chapter 3.4 of Bickel et al. (1998), the projection of $l_{\boldsymbol{\theta}}$ on the orthogonal complement of Υ_λ, i.e., $\Upsilon_\lambda^\perp$, is given by

$$\int [\boldsymbol{g}(\boldsymbol{X}, A) - E\{\boldsymbol{g}(\boldsymbol{X}, A)|T = u, \Delta = 1\}]\, dM(u, \boldsymbol{X}, A).$$

On the other hand, the score function for ψ takes form $\int A\boldsymbol{f}(\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z})dM(u, \boldsymbol{X}, A)$, so

$$\Upsilon_\psi \cap \Upsilon_\lambda^\perp = \Big\{ \int \big[A\boldsymbol{f}(\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z}) - E\{A\boldsymbol{f}(\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z})|T = u, \Delta = 1\}\big]\, dM(u, \boldsymbol{X}, A) : \\ \boldsymbol{f}(u) \text{ is a vector of measurable functions}\Big\}.$$

Therefore, to find the projection of $l_{\boldsymbol{\theta}}$ on the tangent space for (λ, ψ), it suffices to find $\boldsymbol{R}(u)$ that satisfies

$$E\Big(\int \big[\boldsymbol{g}(\boldsymbol{X}, A) - E\{\boldsymbol{g}(\boldsymbol{X}, A)|T = u, \Delta = 1\} \\ - A\boldsymbol{R}(\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z}) + E\{A\boldsymbol{R}(\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z})|T = u, \Delta = 1\}\big] dM(u, \boldsymbol{X}, A) \\ \int \big[A\boldsymbol{f}(\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z}) - E\{A\boldsymbol{f}(\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z})|T = u, \Delta = 1\}\big]\, dM(u, \boldsymbol{X}, A)\Big) = \boldsymbol{0}$$

for any $\boldsymbol{f}(\cdot)$. Equivalently,

$$\int E\Big(\big[\boldsymbol{g}(\boldsymbol{X}, A) - E\{\boldsymbol{g}(\boldsymbol{X}, A)|T = u, \Delta = 1\} \\ - A\boldsymbol{R}(\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z}) + E\{A\boldsymbol{R}(\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z})|T = u, \Delta = 1\}\big] A\boldsymbol{f}(\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z}) \\ \exp\big\{\boldsymbol{\alpha}_0^{\mathrm{T}}\boldsymbol{X} + A\psi_0(\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z})\big\}Y(u)\Big)\lambda_0(u)du = \boldsymbol{0}$$

for all $\boldsymbol{f}(\cdot)$. Thus, $\boldsymbol{R}(\cdot)$ solves the equation

$$E\Big(\big[\boldsymbol{g}(\boldsymbol{X}, A) - E\{\boldsymbol{g}(\boldsymbol{X}, A)|T = u, \Delta = 1\} - A\boldsymbol{R}(\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z}) + E\{A\boldsymbol{R}(\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z})|T = u, \Delta = 1\}\big] \\ A\exp\big\{\boldsymbol{\alpha}_0^{\mathrm{T}}\boldsymbol{X} + A\psi_0(\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z})\big\}Y(u)|\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z}\Big) = \boldsymbol{0},$$

which is equation (3.4). In §S.2 of the Supplementary Materials, we show that the solution to (3.4) exists and is unique. As a result, $\boldsymbol{h}_\psi(u) = \boldsymbol{R}(u)$, and $\boldsymbol{h}_\lambda(t) = E[g(X, A) + A\boldsymbol{R}(\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z})|T = t, \Delta = 1]$. In addition, the efficient score function for $\boldsymbol{\theta}$ as

$$\boldsymbol{S}_{\mathrm{eff},\boldsymbol{\theta}} \equiv \int \boldsymbol{H}(u, A, \boldsymbol{X}, \boldsymbol{Z}) dM(u, A, \boldsymbol{X}, \boldsymbol{Z}).$$

We approximate $\boldsymbol{h}_\psi$ by a B-spline function $\widetilde{\boldsymbol{h}}_\psi$ such that $||\widetilde{\boldsymbol{h}}_\psi - \boldsymbol{h}_\psi||_{W^{1,\infty}(\mathcal{Z})}$ $\leqslant O(K_n^{-2})$ and $||\widetilde{\boldsymbol{h}}_\psi - \boldsymbol{h}_\psi||_{L^2(\mathcal{Z})} \leqslant O(K_n^{-7/2})$. Because

$$\mathbb{P}_n \left\{ \partial l(\boldsymbol{\theta}, \widehat{\boldsymbol{\phi}})/\partial \boldsymbol{\theta}|_{\boldsymbol{\theta}=\widehat{\boldsymbol{\theta}}} \right\} = \mathbf{0}$$

and

$$\mathbb{P}_n \left\{ \partial l(\widehat{\boldsymbol{\theta}}, \widehat{\boldsymbol{\phi}} + \epsilon(\boldsymbol{h}_\lambda, \widetilde{\boldsymbol{h}}_\psi)^{\mathrm{T}})/\partial \epsilon|_{\epsilon=0} \right\} = 0,$$

we have

$$\mathbb{P}_n l_{\boldsymbol{\theta}}(\widehat{\boldsymbol{\theta}}, \widehat{\boldsymbol{\phi}}) - \mathbb{P}_n l_{\boldsymbol{\phi}}(\widehat{\boldsymbol{\theta}}, \widehat{\boldsymbol{\phi}})[(\boldsymbol{h}_\lambda, \widetilde{\boldsymbol{h}}_\psi)^{\mathrm{T}}] = \mathbf{0},$$

or, equivalently,

$$(\mathbb{P}_n - \mathbb{P}) \left\{ (l_{\boldsymbol{\theta}}(\widehat{\boldsymbol{\theta}}, \widehat{\boldsymbol{\phi}}) - l_{\boldsymbol{\phi}}(\widehat{\boldsymbol{\theta}}, \widehat{\boldsymbol{\phi}})[(\boldsymbol{h}_\lambda, \widetilde{\boldsymbol{h}}_\psi)^{\mathrm{T}}] \right\} = -\mathbb{P} \left\{ l_{\boldsymbol{\theta}}(\widehat{\boldsymbol{\theta}}, \widehat{\boldsymbol{\phi}}) - l_{\boldsymbol{\phi}}(\widehat{\boldsymbol{\theta}}, \widehat{\boldsymbol{\phi}})[(\boldsymbol{h}_\lambda, \widetilde{\boldsymbol{h}}_\psi)^{\mathrm{T}}] \right\}. \quad (8)$$

Due to the convergence of $\widehat{\boldsymbol{\theta}}$ and $\widehat{\boldsymbol{\phi}}$ and the Donsker properties of $l_{\boldsymbol{\theta}}(\widehat{\boldsymbol{\theta}}, \widehat{\boldsymbol{\phi}})$ and $l_{\boldsymbol{\phi}}(\widehat{\boldsymbol{\theta}}, \widehat{\boldsymbol{\phi}})$, (8) becomes

$$\begin{aligned}(\mathbb{P}_n - \mathbb{P}) \left\{ (l_{\boldsymbol{\theta}}(\boldsymbol{\theta}_0, \boldsymbol{\phi}_0) - l_{\boldsymbol{\phi}}(\boldsymbol{\theta}_0, \boldsymbol{\phi}_0)[(\boldsymbol{h}_\lambda, \widetilde{\boldsymbol{h}}_\psi)^{\mathrm{T}}] \right\} + o_P(n^{-1/2}) \\ = -\mathbb{P} \left\{ l_{\boldsymbol{\theta}}(\widehat{\boldsymbol{\theta}}, \widehat{\boldsymbol{\phi}}) - l_{\boldsymbol{\phi}}(\widehat{\boldsymbol{\theta}}, \widehat{\boldsymbol{\phi}})[(\boldsymbol{h}_\lambda, \widetilde{\boldsymbol{h}}_\psi)^{\mathrm{T}}] \right\}.\end{aligned}$$

By the first-order Taylor expansion on the right side of the above equation,

$$
\begin{aligned}
&-(1+o_P(1))\mathbb{P}\left\{l_{\boldsymbol{\theta\theta}} - l_{\boldsymbol{\phi\theta}}[(\boldsymbol{h}_\lambda, \boldsymbol{h}_\psi)^{\mathrm{T}}]\right\}(\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0) \\
&\qquad - (1+o_P(1))\mathbb{P}\left\{l_{\boldsymbol{\theta\phi}}[\widehat{\boldsymbol{\phi}} - \boldsymbol{\phi}_0] - l_{\boldsymbol{\phi\phi}}[(\boldsymbol{h}_\lambda, \boldsymbol{h}_\psi)^{\mathrm{T}}, \widehat{\boldsymbol{\phi}} - \boldsymbol{\phi}_0]\right\} \\
&\qquad\qquad + O_p(||\widehat{\boldsymbol{\alpha}} - \boldsymbol{\alpha}_0||^2 + ||\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0||^2 + ||\widehat{\boldsymbol{\phi}} - \boldsymbol{\phi}_0||^2_{L_2(\mathcal{Z})}) + O_P(K_n^{-7}).
\end{aligned}
$$

The choice of $(\boldsymbol{h}_\lambda, \boldsymbol{h}_\psi)^{\mathrm{T}}$ implies that the term $\mathbb{P}\{l_{\boldsymbol{\theta\phi}}[\widehat{\boldsymbol{\phi}} - \boldsymbol{\phi}_0] - l_{\boldsymbol{\phi\phi}}[(\boldsymbol{h}_\lambda, \boldsymbol{h}_\psi)^{\mathrm{T}}, \widehat{\boldsymbol{\phi}} - \boldsymbol{\phi}_0]\}$ is zero. It then follows from the convergence rate of $(\widehat{\boldsymbol{\alpha}}, \widehat{\boldsymbol{\beta}}, \widehat{\phi})$ that

$$
\begin{aligned}
&(\mathbb{P}_n - \mathbb{P})\left\{(l_{\boldsymbol{\theta}}(\boldsymbol{\theta}_0, \boldsymbol{\phi}_0) - l_{\boldsymbol{\phi}}(\boldsymbol{\theta}_0, \boldsymbol{\phi}_0)[(\boldsymbol{h}_\lambda, \widetilde{\boldsymbol{h}}_\psi)^{\mathrm{T}}]\right\} + o_P(n^{-1/2}) \\
&\qquad = -(1+o_P(1))\mathbb{P}\left\{l_{\boldsymbol{\theta\theta}} - l_{\boldsymbol{\phi\theta}}[(\boldsymbol{h}_\lambda, \boldsymbol{h}_\psi)^{\mathrm{T}}]\right\}(\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0) + o_P(n^{-1/2}) + O_P(K_n^{-7}).
\end{aligned}
$$

In §S.3 of the Supplementary Materials, we show that $\mathbb{P}\{l_{\boldsymbol{\theta\theta}} - l_{\boldsymbol{\phi\theta}}[(\boldsymbol{h}_\lambda, \boldsymbol{h}_\psi)^{\mathrm{T}}]\}$ is non-singular. Thus,

$$
\begin{aligned}
&n^{1/2}(\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0) \\
&= -\,\mathbb{P}\left\{l_{\boldsymbol{\theta\theta}} - l_{\boldsymbol{\phi\theta}}[(\boldsymbol{h}_\lambda, \boldsymbol{h}_\psi)^{\mathrm{T}}]\right\}^{-1} n^{1/2}(\mathbb{P}_n - \mathbb{P})\left\{(l_{\boldsymbol{\theta}}(\boldsymbol{\theta}_0, \boldsymbol{\phi}_0) - l_{\boldsymbol{\phi}}(\boldsymbol{\theta}_0, \boldsymbol{\phi}_0)[(\boldsymbol{h}_\lambda, \widetilde{\boldsymbol{h}}_\psi)^{\mathrm{T}}]\right\} + o_P(1) \\
&= -\,\mathbb{P}\left\{l_{\boldsymbol{\theta\theta}} - l_{\boldsymbol{\phi\theta}}[(\boldsymbol{h}_\lambda, \boldsymbol{h}_\psi)^{\mathrm{T}}]\right\}^{-1} n^{1/2}(\mathbb{P}_n - \mathbb{P})\left\{(l_{\boldsymbol{\theta}}(\boldsymbol{\theta}_0, \boldsymbol{\phi}_0) - l_{\boldsymbol{\phi}}(\boldsymbol{\theta}_0, \boldsymbol{\phi}_0)[(\boldsymbol{h}_\lambda, \boldsymbol{h}_\psi)^{\mathrm{T}}]\right\} + o_P(1).
\end{aligned}
$$

The asymptotic normality in Theorem 2 follows. In addition, the limiting covariance matrix is given by $\boldsymbol{\Sigma}^{-1}$, where $\boldsymbol{\Sigma} = E(\boldsymbol{S}_{\mathrm{eff},\boldsymbol{\theta}}\boldsymbol{S}_{\mathrm{eff},\boldsymbol{\theta}}^{\mathrm{T}})$. Clearly, $\boldsymbol{\Sigma}^{-1}$ achieves the semiparametric efficiency bound.

A.1.3 Proof of Theorem 3.3.3

We have established the consistency of $\widehat{\boldsymbol{\alpha}}$, $\widehat{\boldsymbol{\beta}}$, and $\widehat{\psi}$. The consistency for the estimator of the cumulative baseline hazard function can be obtained in a similar manner to Tsiatis (1981). Thus, to prove Theorem 3, it suffices to show that $\widehat{\boldsymbol{R}}(\widehat{\boldsymbol{\beta}}^{\mathrm{T}}\boldsymbol{Z}) \to \boldsymbol{R}(\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z})$ in probability.

Let $\boldsymbol{L}$ and $\widehat{\boldsymbol{L}}$ denote the linear operators in the equations for $\boldsymbol{R}$ and $\widehat{\boldsymbol{R}}$, respectively, and let $\boldsymbol{I}$ denote the identity operator on the same space. Then $(\boldsymbol{I}-\boldsymbol{L})\boldsymbol{R}=\boldsymbol{0}$ and $(\boldsymbol{I}-\widehat{\boldsymbol{L}})\widehat{\boldsymbol{R}}=\boldsymbol{0}$. Since

$$E(A\boldsymbol{R}|T=u,\Delta=1)=E[\exp\{\boldsymbol{\alpha}_0^{\mathrm{T}}\boldsymbol{X}+A\psi_0(\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z})\}Y(u)]^{-1}$$
$$E[(A\boldsymbol{R}\exp\{\boldsymbol{\alpha}_0^{\mathrm{T}}\boldsymbol{X}+A\psi_0(\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z})\}Y(u)],$$

which is uniformly bounded in u, it is clear that $\boldsymbol{L}$ and $\widehat{\boldsymbol{L}}$ are both bounded linear operators mapping the space of bounded variation functions to itself.

For any function with bounded total variation in $\mathcal{Z}$,

$$\begin{aligned}
&||(\boldsymbol{L}-\widehat{\boldsymbol{L}})\boldsymbol{r}(s)||_{L^\infty(\mathcal{Z})} \qquad (9)\\
\leqslant&\Bigg|\Bigg|E\Big[\int A\exp\left\{\boldsymbol{\alpha}_0^{\mathrm{T}}\boldsymbol{Z}+A\psi_0(s)\right\}Y(u)\lambda_0(u)du|\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z}=s\Big]^{-1}\\
&\quad E\Big[\int AY(u)\{\boldsymbol{g}-E(\boldsymbol{g}|T=u,\Delta=1)\}\exp\left\{\boldsymbol{\alpha}_0^{\mathrm{T}}\boldsymbol{X}+A\psi_0(s)\right\}\lambda_0(u)du|\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z}=s\Big]\\
&-\widehat{E}_n\Big[\int A\exp\left\{\widehat{\boldsymbol{\alpha}}^{\mathrm{T}}\boldsymbol{Z}+A\widehat{\psi}(s)\right\}Y(u)d\widehat{\Lambda}(u)|\widehat{\boldsymbol{\beta}}^{\mathrm{T}}\boldsymbol{Z}=s\Big]^{-1}\\
&\quad\widehat{E}_n\Big[\int AY(u)\left\{\widehat{\boldsymbol{g}}-\widetilde{E}_n(\widehat{\boldsymbol{g}}|T=u,\Delta=1)\right\}\exp\left\{\widehat{\boldsymbol{\alpha}}^{\mathrm{T}}\boldsymbol{X}+A\widehat{\psi}(s)\right\}d\widehat{\Lambda}(u)|\widehat{\boldsymbol{\beta}}^{\mathrm{T}}\boldsymbol{Z}=s\Big]\Bigg|\Bigg|_{L^\infty(\mathcal{Z})}\\
&+\Bigg|\Bigg|E\Big[\int A\exp\left\{\boldsymbol{\alpha}_0^{\mathrm{T}}\boldsymbol{X}+A\psi_0(s)\right\}Y(u)\lambda_0(u)du|\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z}=s\Big]^{-1}\\
&\quad E\Big[\int AY(u)E(A\boldsymbol{r}(\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z})|T=u,\Delta=1)\exp\left\{\boldsymbol{\alpha}_0^{\mathrm{T}}\boldsymbol{X}+A\psi_0(s)\right\}\lambda_0(u)du|\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z}=s\Big]\\
&-\widehat{E}_n\Big[\int A\exp\left\{\widehat{\boldsymbol{\alpha}}^{\mathrm{T}}\boldsymbol{X}+A\widehat{\psi}(s)\right\}Y(u)d\widehat{\Lambda}(u)|\widehat{\boldsymbol{\beta}}^{\mathrm{T}}\boldsymbol{Z}=s\Big]^{-1}\\
&\quad\widehat{E}_n\Big[\int AY(u)\widetilde{E}_n(A\boldsymbol{r}(\widehat{\boldsymbol{\beta}}^{\mathrm{T}}\boldsymbol{Z})|T=u,\Delta=1)\exp\left\{\widehat{\boldsymbol{\alpha}}^{\mathrm{T}}\boldsymbol{X}+A\widehat{\psi}(s)\right\}d\widehat{\Lambda}(u)|\widehat{\boldsymbol{\beta}}^{\mathrm{T}}\boldsymbol{Z}=s\Big]\Bigg|\Bigg|_{L^\infty(\mathcal{Z})}.
\end{aligned}$$

Clearly, for any bounded function f, $\widetilde{E}_n(f|T=u,\Delta=1)$ converges to $E(f|T=u,\Delta=1)$ uniformly in u. In addition, with the choice of m_n, $\widehat{E}_n(f|\widehat{\boldsymbol{\beta}}^{\mathrm{T}}\boldsymbol{Z}=s)$ converges to the true conditional expectation of f given $\boldsymbol{\beta}_0^{\mathrm{T}}\boldsymbol{Z}$ by the standard results for histogram-type estimators

(Wasserman, 2006, Chapter 6.2) and the convergence of $\widehat{\boldsymbol{\beta}}$ to $\boldsymbol{\beta}_0$. Thus, the norms of the two terms on the right side of (9) converge to 0, such that $||(\boldsymbol{L}-\widehat{\boldsymbol{L}})\boldsymbol{r}(\cdot)||_{L^\infty(\mathcal{Z})} \to 0$ as n increases.

In §S.3 of the Supplementary Materials, we show that the operator $(\boldsymbol{I}-\boldsymbol{L})$ is invertible. It follows that $(\boldsymbol{I}-\widehat{\boldsymbol{L}})$ is also invertible for large n and moreover,

$$||(\boldsymbol{I}-\boldsymbol{L})^{-1}-(\boldsymbol{I}-\widehat{\boldsymbol{L}})^{-1}|| < c||(\boldsymbol{L}-\widehat{\boldsymbol{L}})|| \to_p 0$$

for some constant c. This gives $||\widehat{\boldsymbol{R}}-\boldsymbol{R}||_{L^\infty(\mathcal{Z})} \to 0$.

A.2 Additional Results for Chapter 3

A.3 Proof for Chapter 4

A.3.1 Invertibility of information operator

Without loss of generality, we assume $\beta_q > 0$. Recall that $\boldsymbol{\theta} = (\boldsymbol{\alpha}^{\mathrm{T}}, \boldsymbol{\beta}_{-q}^{\mathrm{T}})^{\mathrm{T}}$, where $\boldsymbol{\beta}_{-q} = (\beta_1, \cdots, \beta_{q-1})^{\mathrm{T}}$. The log-likelihood from a single subject is

$$\begin{aligned} l(\boldsymbol{\theta},\psi,\Lambda) = {} & \Delta\{\boldsymbol{\alpha}^{\mathrm{T}}\boldsymbol{X} + A\psi(\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z})\} + \Delta\log\lambda(\widetilde{T}) + \\ & \Delta\log G'\left[\exp\left\{\boldsymbol{\alpha}^{\mathrm{T}}\boldsymbol{X} + A\psi(\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z})\right\}\Lambda(\widetilde{T})\right] - G\left[\exp\left\{\boldsymbol{\alpha}^{\mathrm{T}}\boldsymbol{X} + A\psi(\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z})\right\}\Lambda(\widetilde{T})\right] \end{aligned}$$

By differentiating the log-likelihood with respect to $\boldsymbol{\theta}$, ψ and Λ along the submodel $\boldsymbol{\theta}+\epsilon\boldsymbol{d}$, $\psi+\epsilon h_\psi$ and $d\Lambda(1+\epsilon h_\Lambda)$, respectively, we obtain the score operator as

$$\begin{aligned} l_{\boldsymbol{\theta}}(\boldsymbol{\theta},\psi,\Lambda)^{\mathrm{T}}\boldsymbol{d} &= [\Delta + g_{\boldsymbol{\theta},\psi,\Lambda}(\boldsymbol{X},A,\Delta,\widetilde{T})\Lambda(\widetilde{T})] \begin{pmatrix} \boldsymbol{X} \\ A\psi'(\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z})(\boldsymbol{Z}_{-q}-\boldsymbol{\beta}_{-q}Z_q/\beta_q) \end{pmatrix}^{\mathrm{T}} \boldsymbol{d}, \\ l_\psi(\boldsymbol{\theta},\psi,\Lambda)[h_\psi] &= [\Delta + g_{\boldsymbol{\theta},\psi,\Lambda}(\boldsymbol{X},A,\Delta,\widetilde{T})\Lambda(\widetilde{T})]Ah_\psi(\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z}), \\ l_\Lambda(\boldsymbol{\theta},\psi,\Lambda)[h_\Lambda] &= \Delta h_\Lambda(\widetilde{T}) + g_{\boldsymbol{\theta},\psi,\Lambda}(\boldsymbol{X},A,\Delta,\widetilde{T})\int_0^{\widetilde{T}} h(t)dt, \end{aligned}$$

where

$$g_{\boldsymbol{\theta},\psi,\Lambda}(\boldsymbol{X}, A, \Delta, \widetilde{T}) =$$
$$\Big(\Delta G'\left[\exp\left\{\boldsymbol{\alpha}^{\mathrm{T}}\boldsymbol{X} + A\psi(\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z})\right\}\Lambda(\widetilde{T})\right]^{-1} G''\left[\exp\left\{\boldsymbol{\alpha}^{\mathrm{T}}\boldsymbol{X} + A\psi(\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z})\right\}\Lambda(\widetilde{T})\right]$$
$$- G'\left[\exp\left\{\boldsymbol{\alpha}^{\mathrm{T}}\boldsymbol{X} + A\psi(\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z})\right\}\Lambda(\widetilde{T})\right]\Big)\exp\left\{\boldsymbol{\alpha}^{\mathrm{T}}\boldsymbol{X} + A\psi(\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z})\right\}.$$

We write $\boldsymbol{d} = (\boldsymbol{d}_1^{\mathrm{T}}, \boldsymbol{d}_2^{\mathrm{T}})^{\mathrm{T}}$, where $\boldsymbol{d}_1 \in \mathbb{R}^p$, and $\boldsymbol{d}_2 \in \mathbb{R}^{q-1}$. Similarly, $l_{\boldsymbol{\theta}} = (l_{\boldsymbol{\alpha}}^{\mathrm{T}}, l_{\boldsymbol{\beta}_{-q}}^{\mathrm{T}})^{\mathrm{T}}$. In other words, the score operator $S([d, h_\psi, h_\Lambda])$ can be written as $l_{\boldsymbol{\theta}}(\boldsymbol{\theta}, \psi, \Lambda)^{\mathrm{T}}\boldsymbol{d} + l_\psi(\boldsymbol{\theta}, \psi, \Lambda)[h_\psi] + l_\Lambda(\boldsymbol{\theta}, \psi, \Lambda)[h_\Lambda]$.

The dual operator must satisfy

$$< S^*[q(\boldsymbol{X}, A, \Delta, \widetilde{T})], (\boldsymbol{d}, h_\psi, h_\Lambda) >$$
$$= < q(\boldsymbol{X}, A, \Delta, \widetilde{T}), S([\boldsymbol{d}, h_\psi, h_\Lambda]) >$$
$$= E[l_{\boldsymbol{\theta}}(\boldsymbol{\theta}, \psi, \Lambda)^{\mathrm{T}} q]\boldsymbol{d} + E\{l_\psi(\boldsymbol{\theta}, \psi, \Lambda)^{\mathrm{T}}[h_\psi]q\} + E\{l_\Lambda(\boldsymbol{\theta}, \psi, \Lambda)^{\mathrm{T}}[h_\Lambda]q\}$$
$$= E[l_{\boldsymbol{\theta}}(\boldsymbol{\theta}, \psi, \Lambda)^{\mathrm{T}} q(\boldsymbol{X}, A, \Delta, \widetilde{T})]\boldsymbol{d} +$$
$$\int E\{[\Delta + g_{\boldsymbol{\theta},\psi,\Lambda}(\boldsymbol{X}, A, \Delta, \widetilde{T})\Lambda(\widetilde{T})]Ah_\psi(\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z})q(\boldsymbol{X}, A, \Delta, \widetilde{T})|\boldsymbol{X} = x\}f_{\boldsymbol{X}}(x)dx +$$
$$\int E[\Delta q(\boldsymbol{X}, A, \Delta, \widetilde{T})|\widetilde{T} = t]f_{\widetilde{T}}(t)h_\Lambda(t)dt +$$
$$\int E[g_{\boldsymbol{\theta},\psi,\Lambda}(\boldsymbol{X}, A, \Delta, \widetilde{T})q(\boldsymbol{X}, A, \Delta, \widetilde{T})|\widetilde{T} > t]h_\Lambda(t)dt$$

Therefore, $S^*[q] = (E[l_{\boldsymbol{\theta}}(\boldsymbol{\theta}, \psi, \Lambda)^{\mathrm{T}} q], E\{[\Delta + g_{\boldsymbol{\theta},\psi,\Lambda}(\boldsymbol{X}, A, \Delta, \widetilde{T})\Lambda(\widetilde{T})]Ah_\psi(\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z})q|\boldsymbol{X} = x\}$, $E[\Delta q(\boldsymbol{X}, A, \Delta, \widetilde{T})|\widetilde{T} = t]f_{\widetilde{T}}(t) + E[g_{\boldsymbol{\theta},\psi,\Lambda}(\boldsymbol{X}, A, \Delta, \widetilde{T})q(\boldsymbol{X}, A, \Delta, \widetilde{T})|\widetilde{T} > t])^{\mathrm{T}}$. The information operator is

$$\mathcal{L}(\boldsymbol{\theta}, \psi, \Lambda) \equiv S^* S[(\boldsymbol{d}, h_\psi, h_\Lambda)]$$

$$
= \begin{pmatrix} E[l_{\boldsymbol{\theta}} l_{\boldsymbol{\theta}}^{\mathrm{T}}]\boldsymbol{d} \\ E\{l_{\psi}^2[h_{\psi}]|\boldsymbol{X}=x\} \\ E\{l_{\Lambda}[h_{\Lambda}]\Delta|\tilde{T}=t\}f_{\tilde{T}}(t) + E[l_{\Lambda} g_{\boldsymbol{\theta},\psi,\Lambda}|\tilde{T}=t] \end{pmatrix}
$$

$$
+ \begin{pmatrix} E[l_{\boldsymbol{\theta}}(l_{\psi}[h_{\psi}] + l_{\Lambda}[h_{\Lambda}])] \\ E\{l_{\psi}[h_{\psi}](l_{\boldsymbol{\theta}}^{\mathrm{T}}\boldsymbol{d} + l_{\Lambda}[h_{\Lambda}])|\boldsymbol{X}=x\} \\ E\{(l_{\boldsymbol{\theta}}^{\mathrm{T}}\boldsymbol{d} + l_{\psi}[h_{\psi}])\Delta|\tilde{T}=t\}f_{\tilde{T}}(t) + E[(l_{\boldsymbol{\theta}}^{\mathrm{T}}\boldsymbol{d} + l_{\psi}[h_{\psi}])g_{\boldsymbol{\theta},\psi,\Lambda}|\tilde{T}=t] \end{pmatrix}.
$$

The information operator is a Fredholm operator of the first kind, as it is a summation of an invertible operator and an integral operator when $\boldsymbol{\alpha} = \boldsymbol{\alpha}_0, \boldsymbol{\beta} = \boldsymbol{\beta}_{0,-q}, \psi = \psi_0, \Lambda = \Lambda_0$. Furthermore, since h_{ψ} and h_{Λ} are compact, the latter can be shown to be a compact and smooth operator from

$$
\mathcal{H} \equiv \left\{(\boldsymbol{d}, h_{\psi}, h_{\Lambda}) : \boldsymbol{d} \in \mathbb{R}^{p+q-1}, h_{\psi} \in C[0,\tau], h_{\Lambda} \in BV[0,\tau], ||\boldsymbol{d}_2|| \leqslant 1\right\}
$$

to itself, where $BV[0,\tau]$ is the Banach space consisting of all the functions with bounded total variation in $[0,\tau]$ and $C[0,\tau]$ is the Banach space consisting of all the continuous functions in $[0,\tau]$.

We next show that $\mathcal{L}(\boldsymbol{\theta}, \psi, \Lambda)$ is invertible. Following Rudin (1973), it suffices to show that $\mathcal{L}(\boldsymbol{\theta}, \psi, \Lambda)$ is one-to-one. Suppose $\mathcal{L}(\boldsymbol{\theta}_0, \psi_0, \Lambda_0)[\boldsymbol{d}, h_{\psi}, h_{\Lambda}] = 0$; that is, $l_{\boldsymbol{\theta}}^{\mathrm{T}}\boldsymbol{d} + l_{\psi}[h_{\psi}] + l_{\Lambda}[h_{\Lambda}]) = 0$.

First, we let $A = 0$ so that $l_{\boldsymbol{\alpha}}^{\mathrm{T}}\boldsymbol{d}_1 + l_{\Lambda}[h_{\Lambda}] = 0$. This is the score function in the usual transformation model without the single index terms. Thus, from the invertibility of the information in Cox model, we obtain $h_{\Lambda} = 0$ and $\boldsymbol{d}_1 = \boldsymbol{0}$. Now with $l_{\beta_{-q}}^{\mathrm{T}}\boldsymbol{d}_2 + l_{\psi}[h_{\psi}] = 0$, it is necessary that $\psi'(\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z})\boldsymbol{v}^{\mathrm{T}}\boldsymbol{Z} - h_{\psi}(\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z}) = 0$ with $\boldsymbol{v} = \left(\boldsymbol{d}_2^{\mathrm{T}}, -\boldsymbol{d}_2^{\mathrm{T}}\boldsymbol{\beta}_{-q}/\beta_q\right)^{\mathrm{T}}$. This immediately gives $\boldsymbol{v} = c\boldsymbol{\beta}$ and $h_{\psi}(\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z}) = -c\psi'(\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z})\boldsymbol{\beta}^{\mathrm{T}}\boldsymbol{Z}$ for some constant c. Because $\boldsymbol{v}^{\mathrm{T}}\boldsymbol{\beta}_0 = 0$, this is impossible unless $c = 0$. Therefore, $\mathcal{L}(\boldsymbol{\theta}, \psi, \Lambda)$ is invertible.

Finally, the same arguments apply if we consider a different Banach space,

$$\mathcal{H}^* \equiv \{(\boldsymbol{d}, h_\psi, h_\Lambda) : \boldsymbol{d} \in \mathbb{R}^{p+q-1}, h_\psi \in L_2[0,\tau], h_\Lambda \in L_2[0,\tau], ||\boldsymbol{d}_2|| \leqslant 1\}.$$

Therefore, the invertibility of $\mathcal{L}(\boldsymbol{\theta}_0, \psi_0, \Lambda_0)$ implies

$$||\mathcal{L}(\boldsymbol{\theta}_0, \psi_0, \Lambda_0)[\boldsymbol{d}, h_\psi, h_\Lambda]||^2_{L_2(P)} \geqslant c\left\{||\boldsymbol{d}||^2 + ||h_\psi||^2_\infty + ||h_\Lambda||^2_\infty\right\}.$$

Furthermore, we note that $\mathcal{L}(\boldsymbol{\theta}, \psi, \Lambda)$ converges to $\mathcal{L}(\boldsymbol{\theta}_0, \psi_0, \Lambda_0)$ uniformly in the norm $||\boldsymbol{\theta} - \boldsymbol{\theta}_0|| + ||\psi - \psi_0||_\infty + ||\Lambda - \Lambda_0||_\infty$ where $||f||_\infty$ denotes the supreme norm in $[0,\tau]$. We conclude that there exists some ϵ_0 such that whenever $||\boldsymbol{\theta} - \boldsymbol{\theta}_0|| + ||\psi - \psi_0||_\infty + ||\Lambda - \Lambda_0||_\infty < \epsilon_0$, the inequality

$$||\mathcal{L}(\boldsymbol{\theta}, \psi, \Lambda)[\boldsymbol{d}, h_\psi, h_\Lambda]||^2_{L_2(P)} \geqslant c/2\left\{||\boldsymbol{d}||^2 + ||h_\psi||^2_\infty + ||h_\Lambda||^2_\infty\right\} \tag{10}$$

holds.

A.3.2 Proof of Consistency

We will show that there exists a local maximum of the observed data log-likelihood function over the sieve space

$$\begin{aligned}\mathcal{S}_n = \big\{(\boldsymbol{\theta}, \psi, \Lambda) : \boldsymbol{\theta} \in \mathbb{R}^{p+q-1}, \psi(u) = \boldsymbol{\gamma}^{\mathrm{T}}\boldsymbol{B}(u),\\ \text{where } \boldsymbol{B} \text{ is a vector of B-spline bases with knots given in Section 2,}\\ \Lambda \text{ is the step function with jump sizes at the observed events, }\big\}\end{aligned}$$

such that the obtained estimator, $(\widehat{\boldsymbol{\theta}}, \widehat{\psi}, \widehat{\Lambda})$, converges to the true parameters in probability under the norm defined in Theorem 1.

By condition (C.1) and Theorem 6.25 in Schumaker (2007), there exists a function $\widehat{\psi}_0(u) = \boldsymbol{\gamma}_0^{\mathrm{T}}\boldsymbol{B}(u)$ such that $||\widehat{\psi}_0 - \psi_0||_{W^{1,\infty}} = O(K_n^{-2})$ and $||\widehat{\psi}_0 - \psi_0||_{L^2} = O(K_n^{-7/2})$. Then

we consider the following neighborhood of $\widehat{\psi}_0$ in the sieve space

$$\mathcal{N}_{\epsilon_n} = \left\{ \psi(u) = \boldsymbol{\gamma}^{\mathrm{T}} \boldsymbol{B}(u) : \sum_{j=1}^{K_n+1} |\gamma_j - \gamma_{j0}|^2 \leqslant \epsilon_n \right\},$$

where ϵ_n is to be chosen later. For each $\psi \in \mathcal{N}_{\epsilon_n}$, we define

$$(\widehat{\boldsymbol{\theta}}_\psi, \widehat{\Lambda}_\psi) = \mathrm{argmax} P_n l(\boldsymbol{\theta}, \psi, \Lambda),$$

where Λ is a step function with jumps at the observed events.

If we choose ϵ_n so that $K_n^{3/2}\epsilon_n \to 0$, then for $\psi \in \mathcal{N}_{\epsilon_n}$,

$$||\psi - \widehat{\psi}_0||_{BV} \leqslant \sum_{j=1}^{K_n+1} |\gamma_j - \gamma_{j0}| ||B_j'||_\infty = O(K_n)\sqrt{\epsilon_n^2(K_n+1)} \to 0. \tag{11}$$

Therefore, ψ has a bounded total variation. Define

$$\widehat{\Lambda}_0(t) = \int \frac{\sum_{j=1}^n Y_j(t) dN_j(t)}{\sum_{j=1}^n Y_j(t) \exp\{\boldsymbol{\alpha}_0^{\mathrm{T}} \boldsymbol{X}_j + A_j \psi_0(\boldsymbol{\beta}_0^{\mathrm{T}} \boldsymbol{Z}_j)\}}.$$

It is easy to see that $||\widehat{\Lambda}_0 - \Lambda_0||_{BV} = O_p(n^{-1/2})$.

Next, we show that $\mathrm{limsup}_n\{\sup_{\psi \in \mathcal{N}_{\epsilon_n}} \widehat{\Lambda}_\psi(\tau)\}$ is finite with probability tending to one. Define $\zeta_\psi = \widehat{\Lambda}_\psi(\tau)$ and $\bar{\Lambda}_\psi = \widehat{\Lambda}_\psi / \zeta_\psi$ (here all the definition is with the sample size n). Suppose $\zeta_\psi \to \infty$ for some subsequence as n increses to ∞. By the definition of $(\widehat{\boldsymbol{\theta}}_\psi, \widehat{\Lambda}_\psi)$,

$$P_n\{l(\widehat{\boldsymbol{\theta}}_\psi, \psi, \zeta_\psi \bar{\Lambda}_\psi) - l(\widehat{\boldsymbol{\theta}}_\psi, \psi, \bar{\Lambda}_\psi)\} \geqslant 0. \tag{12}$$

By condition (C.1) and (11), $|\widehat{\boldsymbol{\alpha}}_\psi^{\mathrm{T}} \boldsymbol{X} + A\psi(\widehat{\boldsymbol{\beta}}_\psi^{\mathrm{T}} \boldsymbol{Z})|$ is bounded by some $M > 0$. By algebraic manipulation of (12) and the boundedness of $\bar{\Lambda}_\psi$,

$$n^{-1} \sum_{i=1}^n \int_0^{\tau \wedge C_i} log\{\zeta_\psi \sup_{y \leqslant \zeta_\psi e^M} G'(y)\} dN_i(t) - n^{-1} \sum_{i=1}^n I(Y_i(\tau) = 1, C_i \geqslant \tau) G(e^{-M}\zeta_\psi) \geqslant O_p(1).$$

By condition (C.5),

$$\log\{\zeta_\psi \sup_{y \leqslant \zeta_\psi e^M} G'(y)\} \leqslant c_1 G(c_0 \zeta_\psi e^M)$$

for any $c_0 > 0, c_1 > 1$ when n is sufficiently large. For any $\epsilon > 0$, we can choose c_0 such that $G(c_0\zeta_\psi e^M) \leqslant c_1^{-1}\epsilon G(\zeta_\psi e^{-M})$. This implies that

$$\log\{\zeta_\psi \sup_{y \leqslant \zeta_\psi e^M} G'(y)\} \leqslant \epsilon G(\zeta_\psi e^{-M}).$$

Therefore,

$$[n^{-1}\epsilon \sum_{i=1}^{n} N_i(\tau) - n^{-1} \sum_{i=1}^{n} I(Y_i(\tau) = 1, C_i \geqslant \tau)] G(\zeta_n e^{-M}) > -\infty.$$

If we choose an ϵ such that

$$\epsilon \leqslant \frac{Pr(Y(\tau) = 1, C \leqslant \tau)}{2E\{N(\tau)\}},$$

left hand side of the above inequality diverges to $-\infty$, which is a contradiction. Thus, we complete the proof that $\limsup_n \{\sup_{\psi \in \mathcal{N}_{\epsilon_n}} \widehat{\Lambda}_\psi(\tau)\} < \infty$ almost surely.

Since

$$\begin{aligned} P_n l(\boldsymbol{\theta}_\psi, \psi, \widehat{\Lambda}_\psi) \geqslant & P_n l(\boldsymbol{\theta}_0, \psi, \widehat{\Lambda}_0), \\ P_n l_\Lambda(\boldsymbol{\theta}_\psi, \psi, \widehat{\Lambda}_\psi)[h_\Lambda] = & 0, \text{and } P_n l_{\boldsymbol{\theta}}(\boldsymbol{\theta}_\psi, \psi, \widehat{\Lambda}_\psi)^{\mathrm{T}} \boldsymbol{d} = 0, \end{aligned}$$

we obtain

$$(P_n - P) l_\Lambda(\boldsymbol{\theta}_\psi, \psi, \widehat{\Lambda}_\psi)[h_\Lambda] = -P l_\Lambda(\boldsymbol{\theta}_\psi, \psi, \widehat{\Lambda}_\psi)[h_\Lambda]$$

and

$$(P_n - P) l_{\boldsymbol{\theta}}(\boldsymbol{\theta}_\psi, \psi, \widehat{\Lambda}_\psi)^{\mathrm{T}} \boldsymbol{d} = -P l_{\boldsymbol{\theta}}(\boldsymbol{\theta}_\psi, \psi, \widehat{\Lambda}_\psi)^{\mathrm{T}} \boldsymbol{d}.$$

The left-hand sides of the equations are $O_p(n^{-1/2})$ because both l_Λ and $l_{\boldsymbol{\theta}}$ are Donsker due to the fact that both $\widehat{\lambda}_\psi$ and ψ belong to $BV[0, \tau]$. We apply the Taylor expansion at the true

$(\boldsymbol{\theta}_0, \psi_0, \Lambda_0)$ to the right hand side so obtain

$$O_p(n^{-1/2}) = - < \mathcal{L}_{11}(\boldsymbol{\theta}_0, \psi_0, \Lambda_0)[\boldsymbol{d}, h_\Lambda], [\widehat{\boldsymbol{\theta}}_\psi - \boldsymbol{\theta}_0, d\widehat{\Lambda}_\psi - d\Lambda_0] >_{L_2(P)}$$
$$+ o(||\widehat{\boldsymbol{\theta}}_\psi - \boldsymbol{\theta}_0|| + ||\widehat{\Lambda}_\psi - \Lambda_0||_{BV}) + O_p(||\psi - \psi_0||_{L_2}),$$

where $\mathcal{L}_{11}$ is the operator in $\mathcal{L}$ corresponding to $\boldsymbol{\theta}$ and Λ. Using the invertibility of $\mathcal{L}_{11}$, we have

$$||\widehat{\Lambda}_\psi - \Lambda_0||_{BV} + ||\widehat{\boldsymbol{\theta}}_\psi - \boldsymbol{\theta}_0|| = A_n(n^{-1/2} + ||\widehat{\psi} - \psi_0||_{L_2}), \tag{13}$$

where $\sup_{\psi \in \mathcal{N}_{\epsilon_n}} |A_n|$ is a bounded random variable.

We now consider

$$B_n \equiv P_n l(\widehat{\boldsymbol{\theta}}_\psi, \psi, \widehat{\Lambda}_\psi) - P_n l(\boldsymbol{\theta}_0, \widehat{\psi}_0, \widehat{\Lambda}_0).$$

First,

$$B_n = (P_n - P)\left\{l(\widehat{\boldsymbol{\theta}}_\psi, \psi, \widehat{\Lambda}_\psi) - l(\boldsymbol{\theta}_0, \widehat{\psi}_0, \widehat{\Lambda}_0)\right\} + P\left\{l(\widehat{\boldsymbol{\theta}}_\psi, \psi, \widehat{\Lambda}_\psi) - l(\boldsymbol{\theta}_0, \widehat{\psi}_0, \widehat{\Lambda}_0)\right\}.$$

The first term on the right hand side is equal to $C_n n^{-1/2}$ where $\sup_{\psi \in \mathcal{N}_{\epsilon_n}} |C_n| \to 0$ in probability. For the second term, we apply the expansion at the true values. The first order in the expansion vanishes and the second order in the expansion is

$$-\left\langle \mathcal{L}(\boldsymbol{\theta}^*, \psi^*, \Lambda^*)(\widehat{\boldsymbol{\theta}}_\psi - \boldsymbol{\theta}_0, \psi - \psi_0, d\widehat{\Lambda}_\psi/\widehat{\Lambda}_0 - \lambda_0), (\widehat{\boldsymbol{\theta}}_\psi - \boldsymbol{\theta}_0, \psi - \psi_0, d\widehat{\Lambda}_\psi/\widehat{\Lambda}_0 - \lambda_0)\right\rangle_{L_2(P)}$$
$$+ O(||\widehat{\Lambda}_0||_\infty^2 + ||\widehat{\psi}_0 - \psi_0||_\infty^2),$$

where $(\boldsymbol{\theta}^*, \psi^*, \Lambda^*)$ is between $(\widehat{\boldsymbol{\theta}}_\psi, \psi, \widehat{\Lambda}_\psi)$ and $(\boldsymbol{\theta}_0, \psi_0, \Lambda_0)$. Using the result in (10), we obtain

$$B_n = C_n n^{-1/2} - c_1/2||\psi - \psi_0||_{L_2}^2 + D_n(n^{-1} + K_n^{-7}).$$

Therefore, if $\psi \in \delta\mathcal{N}_{\epsilon_n}$, the results from de Boor (1978) gives $||\psi - \psi_0||_{L_2}^2 \geqslant c_2\epsilon_n^2$ so that

$$B_n \leqslant \sup_{\psi\in\mathcal{N}_{\epsilon_n}}\{C_n n^{-1/2} + D_n(n^{-1} + K_n^{-7})\} - c_1c_2\epsilon_n^2/2.$$

Consequently, if we choose

$$\epsilon_n^2 = 4c_1^{-1}c_2^{-1}\sup_{\psi\in\mathcal{N}_{\epsilon_n}}\{C_n n^{-1/2} + D_n(n^{-1} + K_n^{-7})\},$$

then $B_n < 0$. Hence, there exists a local maximum $\widehat{\psi}$ within this neighborhood. Furthermore, $||\widehat{\psi} - \psi_0||_{BV} \to 0$ by (11) and

$$||\widehat{\psi} - \psi_0||_{L_2}^2 \leqslant ||\widehat{\psi} - \widehat{\psi}_0||_{L_2}^2 + O(K_n^{-7}) \leqslant \epsilon_n^2 + K_n^{-7} = o_p(n^{-1/2})$$

according to Condition (C4). By (13), the corresponding $(\widehat{\boldsymbol{\theta}}_{\widehat{\psi}}, \widehat{\Lambda}_{\widehat{\psi}})$ satisfies

$$||\widehat{\Lambda} - \Lambda_0||_{BV} + ||\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0|| = O_p(n^{-1/2}) + ||\widehat{\psi} - \psi_0||_{L_2} = o_p(n^{-1/4}).$$

A.3.3 Proof of Asymptotic Distribution

The least favorable direction h_λ and h_ψ can be constructed from the expression of the dual operator and information operator in section A.2.1. Write $\boldsymbol{\phi} = (\lambda, \psi)^{\mathrm{T}}$. Because

$$\mathbb{P}_n\left\{\partial l(\boldsymbol{\theta}, \widehat{\boldsymbol{\psi}})/\partial\boldsymbol{\theta}|_{\boldsymbol{\theta}=\widehat{\boldsymbol{\theta}}}\right\} = \mathbf{0}$$

and

$$\mathbb{P}_n\big\{\partial l(\widehat{\boldsymbol{\theta}}, \widehat{\boldsymbol{\phi}} + \epsilon(\boldsymbol{h}_\lambda, \widetilde{\boldsymbol{h}}_\psi)^{\mathrm{T}})/\partial\epsilon|_{\epsilon=0}\big\} = 0,$$

we have

$$\mathbb{P}_n l_{\boldsymbol{\theta}}(\widehat{\boldsymbol{\theta}}, \widehat{\boldsymbol{\phi}}) - \mathbb{P}_n l_{\boldsymbol{\phi}}(\widehat{\boldsymbol{\theta}}, \widehat{\boldsymbol{\phi}})[(\boldsymbol{h}_\lambda, \widetilde{\boldsymbol{h}}_\psi)^{\mathrm{T}}] = \mathbf{0},$$

or, equivalently,

$$(\mathbb{P}_n - \mathbb{P})\left\{(l_{\boldsymbol{\theta}}(\widehat{\boldsymbol{\theta}}, \widehat{\boldsymbol{\phi}}) - l_{\boldsymbol{\phi}}(\widehat{\boldsymbol{\theta}}, \widehat{\boldsymbol{\phi}})[(\boldsymbol{h}_\lambda, \widetilde{\boldsymbol{h}}_\psi)^{\mathrm{T}}]\right\} = \\ - \mathbb{P}\left\{l_{\boldsymbol{\theta}}(\widehat{\boldsymbol{\theta}}, \widehat{\boldsymbol{\phi}}) - l_{\boldsymbol{\phi}}(\widehat{\boldsymbol{\theta}}, \widehat{\boldsymbol{\phi}})[(\boldsymbol{h}_\lambda, \widetilde{\boldsymbol{h}}_\psi)^{\mathrm{T}}]\right\}. \quad (14)$$

Due to the convergence of $\widehat{\boldsymbol{\theta}}$ and $\widehat{\boldsymbol{\phi}}$, we could verify that $l_{\boldsymbol{\theta}}(\widehat{\boldsymbol{\theta}}, \widehat{\boldsymbol{\phi}})$ and $l_{\boldsymbol{\phi}}(\widehat{\boldsymbol{\theta}}, \widehat{\boldsymbol{\phi}})$ are P-Donsker, because $\Lambda(t)$ is increasing in t, and the other terms are bounded by our assumptions. Therefore, (14) becomes

$$(\mathbb{P}_n - \mathbb{P})\left\{(l_{\boldsymbol{\theta}}(\boldsymbol{\theta}_0, \boldsymbol{\phi}_0) - l_{\boldsymbol{\phi}}(\boldsymbol{\theta}_0, \boldsymbol{\phi}_0)[(\boldsymbol{h}_\lambda, \widetilde{\boldsymbol{h}}_\psi)^{\mathrm{T}}]\right\} + o_P(n^{-1/2}) \\ = -\mathbb{P}\left\{l_{\boldsymbol{\theta}}(\widehat{\boldsymbol{\theta}}, \widehat{\boldsymbol{\phi}}) - l_{\boldsymbol{\phi}}(\widehat{\boldsymbol{\theta}}, \widehat{\boldsymbol{\phi}})[(\boldsymbol{h}_\lambda, \widetilde{\boldsymbol{h}}_\psi)^{\mathrm{T}}]\right\}.$$

By the first-order Taylor expansion on the right side of the above equation,

$$-(1 + o_P(1))\mathbb{P}\left\{l_{\boldsymbol{\theta}\boldsymbol{\theta}} - l_{\boldsymbol{\phi}\boldsymbol{\theta}}[(\boldsymbol{h}_\lambda, \boldsymbol{h}_\psi)^{\mathrm{T}}]\right\}(\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0) \\ - (1 + o_P(1))\mathbb{P}\left\{l_{\boldsymbol{\theta}\boldsymbol{\phi}}[\widehat{\boldsymbol{\phi}} - \boldsymbol{\phi}_0] - l_{\boldsymbol{\phi}\boldsymbol{\phi}}[(\boldsymbol{h}_\lambda, \boldsymbol{h}_\psi)^{\mathrm{T}}, \widehat{\boldsymbol{\phi}} - \boldsymbol{\phi}_0]\right\} \\ + O_p(||\widehat{\boldsymbol{\alpha}} - \boldsymbol{\alpha}_0||^2 + ||\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0||^2 + ||\widehat{\boldsymbol{\phi}} - \boldsymbol{\phi}_0||^2_{L_2(\mathcal{Z})}) + O_P(K_n^{-7}).$$

The choice of $(\boldsymbol{h}_\lambda, \boldsymbol{h}_\psi)^{\mathrm{T}}$ implies that the term $\mathbb{P}\{l_{\boldsymbol{\theta}\boldsymbol{\phi}}[\widehat{\boldsymbol{\phi}} - \boldsymbol{\phi}_0] - l_{\boldsymbol{\phi}\boldsymbol{\phi}}[(\boldsymbol{h}_\lambda, \boldsymbol{h}_\psi)^{\mathrm{T}}, \widehat{\boldsymbol{\phi}} - \boldsymbol{\phi}_0]\}$ is zero. It then follows from the convergence rate of $(\widehat{\boldsymbol{\alpha}}, \widehat{\boldsymbol{\beta}}, \widehat{\phi})$ that

$$(\mathbb{P}_n - \mathbb{P})\left\{(l_{\boldsymbol{\theta}}(\boldsymbol{\theta}_0, \boldsymbol{\phi}_0) - l_{\boldsymbol{\phi}}(\boldsymbol{\theta}_0, \boldsymbol{\phi}_0)[(\boldsymbol{h}_\lambda, \widetilde{\boldsymbol{h}}_\psi)^{\mathrm{T}}]\right\} + o_P(n^{-1/2}) \\ = -(1 + o_P(1))\mathbb{P}\left\{l_{\boldsymbol{\theta}\boldsymbol{\theta}} - l_{\boldsymbol{\phi}\boldsymbol{\theta}}[(\boldsymbol{h}_\lambda, \boldsymbol{h}_\psi)^{\mathrm{T}}]\right\}(\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0) + o_P(n^{-1/2}) + O_P(K_n^{-7}).$$

Since we showed $\mathbb{P}\{l_{\boldsymbol{\theta\theta}} - l_{\boldsymbol{\phi\theta}}[(\boldsymbol{h}_\lambda, \boldsymbol{h}_\psi)^{\mathrm{T}}]\}$ is non-singular in subsection A.2.1 Thus,

$$
\begin{aligned}
& n^{1/2}(\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0) \\
= & -\mathbb{P}\left\{l_{\boldsymbol{\theta\theta}} - l_{\boldsymbol{\phi\theta}}[(\boldsymbol{h}_\lambda, \boldsymbol{h}_\psi)^{\mathrm{T}}]\right\}^{-1} \\
& n^{1/2}(\mathbb{P}_n - \mathbb{P})\left\{(l_{\boldsymbol{\theta}}(\boldsymbol{\theta}_0, \boldsymbol{\phi}_0) - l_{\boldsymbol{\phi}}(\boldsymbol{\theta}_0, \boldsymbol{\phi}_0)[(\boldsymbol{h}_\lambda, \widetilde{\boldsymbol{h}}_\psi)^{\mathrm{T}}]\right\} + o_P(1) \\
= & -\mathbb{P}\left\{l_{\boldsymbol{\theta\theta}} - l_{\boldsymbol{\phi\theta}}[(\boldsymbol{h}_\lambda, \boldsymbol{h}_\psi)^{\mathrm{T}}]\right\}^{-1} \\
& n^{1/2}(\mathbb{P}_n - \mathbb{P})\left\{(l_{\boldsymbol{\theta}}(\boldsymbol{\theta}_0, \boldsymbol{\phi}_0) - l_{\boldsymbol{\phi}}(\boldsymbol{\theta}_0, \boldsymbol{\phi}_0)[(\boldsymbol{h}_\lambda, \boldsymbol{h}_\psi)^{\mathrm{T}}]\right\} + o_P(1).
\end{aligned}
$$

The asymptotic normality in Theorem 4.2 follows. In addition, the limiting covariance matrix is given by $\boldsymbol{\Sigma}^{-1}$, where $\boldsymbol{\Sigma} = E(\boldsymbol{S}_{\mathrm{eff},\boldsymbol{\theta}}\boldsymbol{S}_{\mathrm{eff},\boldsymbol{\theta}}^{\mathrm{T}})$. Clearly, $\boldsymbol{\Sigma}^{-1}$ achieves the semiparametric efficiency bound.

A.4 Proof for Chapter 5

The proof of Theorems 5.3.1–5.3.3 closely follow the steps in Liu and Zeng (2013). We sketch the outline in this section. First, we establish the following two lemma under conditions (C1)–(C5). These lemmas are straightforward implications from the proof in section A.2

Lemma A.4.1. *Let $U_n(\boldsymbol{\beta}) = \partial l_n(\boldsymbol{\beta})/\boldsymbol{\beta}$, where l_n is the log-likelihood with the transformation single-index model. Then $n^{-1/2}U_n(\boldsymbol{\beta}) = O_p(1)$.*

Lemma A.4.2. *Let $V_n(\boldsymbol{\beta}) = -\partial U_n(\boldsymbol{\beta})/\boldsymbol{\beta}^{\mathrm{T}}$. Then uniformly in $\boldsymbol{\beta}$, $n^{-1}V_n(\boldsymbol{\beta}) \to V(\boldsymbol{\beta})$ almost surely, which does not depend on the data. The matrix $V(\boldsymbol{\beta})$ is positive definite; therefore, $l_n(\boldsymbol{\beta})$ is a concave function when n is sufficiently large.*

A.4.1 Consistency of the adaptive lasso estimator

The penalized objective function

$$
Q_n(\boldsymbol{\beta}) = l_n(\boldsymbol{\beta}) - n\rho_n \sum_{j=1}^{q} |\widetilde{\beta}_j|^{-1}|\beta_j|.
$$

is convex by lemma (A.3.2) and the convexity of the penalty functions when n is large enough. Therefore, there exists a unique maximum of $Q_n(\boldsymbol{\beta})$ for sufficiently large n. It suffices to

show that there exists a local maximum in the ball around the true value $\boldsymbol{\beta}_0$ with probability 1. In other words, we aim to show that for any given $\epsilon > 0$, there exists a constant $C > 0$ such that

$$P\left\{\sup_{||\boldsymbol{u}||\leqslant C} Q_n(\boldsymbol{\beta}_0 + n^{-1/2}\boldsymbol{u}) < Q_n(\boldsymbol{\beta}_0)\right\} \geqslant 1 - \epsilon.$$

Note that

$$n^{-1}\{Q_n(\boldsymbol{\beta}_0 + n^{-1/2}\boldsymbol{u}) - Q_n(\boldsymbol{\beta}_0)\} \leqslant$$
$$n^{-1}\{l_n(\boldsymbol{\beta}_0 + n^{-1/2}\boldsymbol{u}) - l_n(\boldsymbol{\beta}_0)\} - n^{-1/2}\rho_n \sum_{j=1}^{q} |\widetilde{\beta}_j|^{-1}|u_j| \quad (15)$$

By the second order Taylor expansion, the first term in (15) is

$$\begin{aligned} & n^{-1}(n^{-1/2}\boldsymbol{u})^{\mathrm{T}} U_n(\boldsymbol{\beta}_0) - (2n)^{-1}(n^{-1/2}\boldsymbol{u})^{\mathrm{T}} V_n(\boldsymbol{\beta}^*)(n^{-1/2}\boldsymbol{u}) \\ =& n^{-1}O_p(1)\sum_{j=1}^{q}|u_j| - (2n)^{-1}\boldsymbol{u}^{\mathrm{T}}\{V(\boldsymbol{\beta}_0) + o_p(1)\}\boldsymbol{u}, \end{aligned}$$

where $\boldsymbol{\beta}^*$ is between $\boldsymbol{\beta}_0$ and $\boldsymbol{\beta}_0 + n^{-1/2}\boldsymbol{u}$. By the first order Taylor expansion, the second term in (15) is

$$\begin{aligned} & n^{-1/2}\rho_n \sum_{j=1}^{q}|u_j|\left\{|\beta_{j0}|^{-1} - |\beta_{j0}|^{-2} sign(\beta_{j0})(\widetilde{\beta}_j - \beta_{j0}) + o_p(|\widetilde{\beta}_j - \beta_{j0}|)\right\} \\ =& n^{-1/2}\rho_n \sum_{j=1}^{q}\{|\beta_{j0}|^{-1} + O_p(n^{-1/2})\}|u_j| \leqslant n^{-1}O_p(1)\sum_{j=1}^{q}|u_j|, \end{aligned}$$

where the equality follows from the convergence rate of $\widetilde{\boldsymbol{\beta}}$ and the inequality follows from $n^{1/2}\rho_n = O_p(1)$. That is,

$$\begin{aligned} & n^{-1}\{Q_n(\boldsymbol{\beta}_0 + n^{-1/2}\boldsymbol{u}) - Q_n(\boldsymbol{\beta}_0)\} \\ & \leqslant -(2n)^{-1}\boldsymbol{u}^{\mathrm{T}}\{V(\boldsymbol{\beta}_0) + o_p(1)\}\boldsymbol{u} + n^{-1}O_p(1)\sum_{j=1}^{q}|u_j| - n^{-1}O_p(1)\sum_{j=1}^{q}|u_j|. \end{aligned}$$

In the above inequality, the first term on the right hand side is of second order of $\boldsymbol{u}$, while all the other terms are of the first order. Therefore, by choosing a constant C large enough, the first term dominates the rest of the expression. By lemma A.3.2, $n^{-1}\{Q_n(\boldsymbol{\beta}_0 + n^{-1/2}\boldsymbol{u}) - Q_n(\boldsymbol{\beta}_0)\} < 0$ and the proof is completed.

A.4.2 Oracle property of the adaptive lasso estimator

For β_j in $\boldsymbol{\beta}_1 = (\beta_1, \ldots, \beta_{q_1})^{\mathrm{T}}$, we have

$$0 = \partial Q_n(\boldsymbol{\beta})/\partial\beta_j|_{\boldsymbol{\beta}=\widehat{\boldsymbol{\beta}}} = n^{1/2}\left\{n^{-1/2}\partial Q_n(\boldsymbol{\beta})/\partial\beta_j|_{\boldsymbol{\beta}=\widehat{\boldsymbol{\beta}}} - n^{1/2}\rho_n|\widetilde{\beta}_j|^{-1}sign(\widehat{\beta}_j)\right\}$$

By the first order Taylor expansion and lemmas A.3.1–A.3.2, we have

$$\begin{aligned}
0 =& n^{1/2}\left\{n^{-1/2}U_{jn}(\boldsymbol{\beta}_0) + n^{-1}V_{jjn}(\boldsymbol{\beta}^*)n^{1/2}(\widehat{\beta}_j - \beta_{j0}) - n\rho_n|n^{1/2}\widetilde{\beta}_j|^{-1}sign(\widehat{\beta}_j)\right\} \\
=& n^{1/2}\left\{O_p(1) + V_{jj}(\boldsymbol{\beta}_0)n^{1/2}(\widehat{\beta}_j - \beta_{j0}) - n\rho_n|n^{1/2}\widetilde{\beta}_j|^{-1}sign(\widehat{\beta}_j)\right\} \\
=& n^{1/2}\left\{O_p(1) - n^{1/2}O_p(1)sign(\widehat{\beta}_j)\right\}.
\end{aligned}$$

In the above derivation, $U_{jn}(\boldsymbol{\beta}_0)$ is the jth element of $U_n(\boldsymbol{\beta}_0)$, $V_{jjn}(\boldsymbol{\beta}^*)$ is the (j,j)th element of $V_n(\boldsymbol{\beta}^*)$, $V_{jj}(\boldsymbol{\beta}_0)$ is the (j,j)th element of $V(\boldsymbol{\beta}_0)$, and $\boldsymbol{\beta}^*$ is between $\boldsymbol{\beta}_0$ and $\widehat{\boldsymbol{\beta}}$. The last line follows from the convergence rate of $\widehat{\boldsymbol{\beta}}$ and $\widetilde{\boldsymbol{\beta}}$. As n goes to infinity, $\widehat{\boldsymbol{\beta}}_1 = 0$ with probability tending to 0.

A.4.3 Asymptotic distribution of the adaptive lasso estimator

Since section A.3.2 proves $P(\widehat{\boldsymbol{\beta}}_{1n} = 0) \to 1$, we only need to derive the asymptotic expansion of $\widehat{\boldsymbol{\beta}}_{2n}$ in the probability set $\{\widehat{\boldsymbol{\beta}}_{1n}\} = 0$. Let $U_{2n}(\boldsymbol{\beta})$ denotes the last q_2 elements of $U(\boldsymbol{\beta})$ and V_{22n} the lower $q_2 \times q_2$ submatrix of $V_n(\boldsymbol{\beta})$. Then

$$\begin{aligned}
0 =& \partial Q_n(\boldsymbol{\beta})/\partial\boldsymbol{\beta}_2|_{\boldsymbol{\beta}=(\mathbf{0},\widehat{\boldsymbol{\beta}}_{2n})^{\mathrm{T}}} \\
=& \partial l_n(\boldsymbol{\beta})/\partial\boldsymbol{\beta}_2|_{\boldsymbol{\beta}=(\mathbf{0},\widehat{\boldsymbol{\beta}}_{2n})^{\mathrm{T}}} - n\rho_n\left(|\widetilde{\beta}_{q_1+1}|^{-1}sign(\widehat{\beta}_{q_1+1}), \ldots, |\widetilde{\beta}_{q_2}|^{-1}sign(\widehat{\beta}_{q_2})\right)^{\mathrm{T}} \\
=& U_{2n}(\boldsymbol{\beta}_0) - V_{22n}(\boldsymbol{\beta}^*)(\widehat{\boldsymbol{\beta}}_{2n} - \boldsymbol{\beta}_{20}) - n\rho_n\left(|\widetilde{\beta}_{q_1+1}|^{-1}sign(\widehat{\beta}_{q_1+1}), \ldots, |\widetilde{\beta}_{q_2}|^{-1}sign(\widehat{\beta}_{q_2})\right)^{\mathrm{T}},
\end{aligned}$$

where $\boldsymbol{\beta}^*$ is between $\boldsymbol{\beta}_0$ and $\widetilde{\boldsymbol{\beta}}$. By Lemma A.3.2, the convergence of $\widetilde{\boldsymbol{\beta}}$ and $\widehat{\boldsymbol{\beta}}$ and the rate of ρ_n, we have

$$
\begin{aligned}
& n^{1/2}(\widehat{\boldsymbol{\beta}}_{2n} - \boldsymbol{\beta}_{20}) \\
=& \{n^{-1}V_{22n}(\boldsymbol{\beta}^*)\}^{-1}\left[n^{-1/2}U_{2n}(\boldsymbol{\beta}_0) - n^{1/2}\rho_n\left(|\widetilde{\beta}_{q_1+1}|^{-1}sign(\widehat{\beta}_{q_1+1}), \ldots, |\widetilde{\beta}_{q_2}|^{-1}sign(\widehat{\beta}_{q_2})\right)^{\mathrm{T}}\right] \\
=& V_{22}(\boldsymbol{\beta}_0)^{-1}\{n^{-1/2}U_{2n}(\boldsymbol{\beta}_0)\} + o_p(1)
\end{aligned}
$$

As the expression above only involves the derivatives concerning the derivatives of the log-likelihood but not the penalty term, the asymptotic distribution of $\boldsymbol{\beta}_{2n}$ can be immediately obtained from sections A.2.1-A.2.3. With the same arguments, the asymptotic variance attains the semiparametric efficiency bound.

A.4.4 Asymptotic distribution of the post-selection maximum likelihood estimator

The proof for the asymptotic distribution of the post-selection maximum likelihood estimator is from sections A.2.1-A.2.3.

REFERENCES

Akaike, H. (1970). Statistical predictor identification. *Annals of the Institute of Statistical Mathematics* **22,** 203–217.

Audibert, J.-Y., Tsybakov, A. B., et al. (2007). Fast learning rates for plug-in classifier. *The Annals of Statistics* **35,** 608–633.

Bai, X., Tsiatis, A. A., Lu, W., and Song, R. (2017a). Optimal treatment regimes for survival endpoints using a locally-efficient doubly-robust estimator from a classification perspective. *Lifetime data analysis* **23,** 585–604.

Bai, X., Tsiatis, A. A., Lu, W., and Song, R. (2017b). Optimal treatment regimes for survival endpoints using a locally-efficient doubly-robust estimator from a classification perspective. *Lifetime data analysis* **23,** 585–604.

Behncke, S., Frölich, M., and Lechner, M. (2009). Targeting labour market programmesâresults from a randomized experiment. *Swiss Journal of Economics and Statistics* **145,** 221–268.

Berwin, A. T. and Weingessel, A. (2013). *quadprog: Functions to solve Quadratic Programming Problems.* R package, version 1.5-5.

Bezdek, J. C. and Hathaway, R. J. (2003). Convergence of alternating optimization. *Neural, Parallel and Scientific Computations* **11,** 351–368.

Bickel, P. J., Klaassen, C. A., Ritov, Y., and Wellner, J. A. (1998). *Efficient and Adaptive Estimation for Semiparametric Models.* Springer New York.

Breiman, L. (2001). Random forests. *Machine Learning* **45,** 5–32.

Breslow, N. E. (1972). Discussion of the paper by D. R. Cox. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **34,** 216–217.

Brillinger, D. R. (1982). A generalized linear model with âgaussianâ regressor variables. In *Selected Works of David Brillinger*, pages 589–606. Springer.

Candes, E., Tao, T., et al. (2007). The dantzig selector: Statistical estimation when p is much larger than n. *The annals of Statistics* **35,** 2313–2351.

Carroll, R. J., Fan, J., Gijbels, I., and Wand, M. P. (1997). Generalized partially linear single-index models. *Journal of the American Statistical Association* **92,** 477–489.

Castiglione-Gertsch, M., Price, K., Goldhirsch, A., Coates, A., Colleoni, M., Nasi, M., Bernhard, J., Zahrich, D., Bonetti, M., and Gelber, R. (2002). Endocrine responsiveness and tailoring adjuvant therapy for postmenopausal lymph node–negative breast cancer: A randomized trial. *Journal of the National Cancer Institute* **94,** 1054–1065.

Chatfield, C. (1995). Model uncertainty, data mining and statistical inference. *Journal of the Royal Statistical Society: Series A (Statistics in Society)* **158,** 419–444.

Chiou, J.-M. and Müller, H.-G. (1998). Quasi-likelihood regression with unknown link and variance functions. *Journal of the American Statistical Association* **93,** 1376–1387.

Chiou, J.-M., Müller, H.-G., et al. (1999). Nonparametric quasi-likelihood. *The Annals of Statistics* **27,** 36–64.

Cox, D. R. (1972). Regression models and life tables (with discussion). *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **34,** 187–220.

Dabrowska, D. M. and Doksum, K. A. (1988). Partial likelihood in transformation models with censored data. *Scandinavian journal of statistics* pages 1–23.

De Boor, C. (1978). *A Practical Guide to Splines*, volume 27. Springer-Verlag New York.

Diaz, I., Savenkov, O., and Ballman, K. (2018). Targeted learning ensembles for optimal individualized treatment rules with time-to-event outcomes. *Biometrika* **105,** 723–738.

Duan, N. and Li, K.-C. (1991). Slicing regression: a link-free regression method. *The Annals of Statistics* pages 505–530.

Fan, A. et al. (2016). New statistical methods for precision medicine: Variable selection for optimal dynamic treatment regimes and subgroup detection.

Fan, J. (1993). Local linear regression smoothers and their minimax efficiencies. *The Annals of Statistics* **21,** 196–216.

Fan, J., Gijbels, I., King, M., et al. (1997). Local likelihood and local partial likelihood in hazard regression. *The Annals of Statistics* **25,** 1661–1690.

Fan, J. and Li, R. (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American statistical Association* **96,** 1348–1360.

Gasser, T., Muller, H.-G., and Mammitzsch, V. (1985). Kernels for nonparametric curve estimation. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **47,** 238–252.

Geng, Y., Zhang, H. H., and Lu, W. (2015). On optimal treatment regimes selection for mean survival time. *Statistics in Medicine* **34,** 1169–1184.

Goldberg, Y. and Kosorok, M. R. (2012). Q-learning with censored data. *Annals of Statistics* **40,** 529.

Group, I. B. C. S. et al. (2002). Endocrine responsiveness and tailoring adjuvant therapy for postmenopausal lymph node–negative breast cancer: A randomized trial. *JNCI: Journal of the National Cancer Institute* **94,** 1054–1065.

Gunter, L., Zhu, J., and Murphy, S. (2011). Variable selection for qualitative interactions in personalized medicine while controlling the family-wise error rate. *Journal of biopharmaceutical statistics* **21,** 1063–1078.

Hamburg, M. A. and Collins, F. S. (2010). The path to personalized medicine. *New England Journal of Medicine* **2010,** 301–304.

Hammer, S. M., Katzenstein, D. A., Hughes, M. D., Gundacker, H., Schooley, R. T., Haubrich, R. H., Henry, W. K., Lederman, M. M., Phair, J. P., Niu, M., et al. (1996). A trial comparing nucleoside monotherapy with combination therapy in HIV-infected adults with CD4 cell counts from 200 to 500 per cubic millimeter. *New England Journal of Medicine* **335,** 1081–1090.

Hardle, W., Hall, P., Ichimura, H., et al. (1993). Optimal smoothing in single-index models. *The annals of Statistics* **21,** 157–178.

Harrell, F. E., Lee, K. L., and Mark, D. B. (1996). Multivariable prognostic models: issues in developing models, evaluating assumptions and adequacy, and measuring and reducing errors. *Statistics in medicine* **15,** 361–387.

He, X. and Shi, P. (1998). Monotone B-spline smoothing. *Journal of the American Statistical Association* **93,** 643–650.

Huang, J. et al. (1999). Efficient estimation of the partly linear additive cox model. *The Annals of Statistics* **27,** 1536–1563.

Huang, J. Z. and Liu, L. (2006). Polynomial spline estimation and inference of proportional hazards regression models with flexible relative risk form. *Biometrics* **62,** 793–802.

Ichimura, H. (1993). Semiparametric least squares (sls) and weighted sls estimation of single-index models. *Journal of Econometrics* **58,** 71–120.

Imai, K., Ratkovic, M., et al. (2013). Estimating treatment effect heterogeneity in randomized program evaluation. *The Annals of Applied Statistics* **7,** 443–470.

Jiang, R., Lu, W., Song, R., and Davidian, M. (2017). On estimation of optimal treatment regimes for maximizing t-year survival probability. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **79,** 1165–1185.

Jiang, R., Lu, W., Song, R., Hudgens, M. G., and Naprvavnik, S. (2017). Doubly robust estimation of optimal treatment regimes for survival data-with application to an HIV/AIDS study. *The Annals of Applied Statistics* **11,** 1763–1786.

Kalbfleisch, J. D. and Prentice, R. L. (1980). *The statistical analysis of failure time data*, volume 360. John Wiley & Sons.

Kang, S., Lu, W., and Zhang, J. (2018). On estimation of the optimal treatment regime with the additive hazards model. *Statistica Sinica* **28,** 1539.

Laurie, J. A., Moertel, C., Fleming, T., Wieand, H., Leigh, J., Rubin, J., McCormack, G., Gerstner, J., Krook, J., and Malliard, J. (1989). Surgical adjuvant therapy of large-bowel carcinoma: an evaluation of levamisole and the combination of levamisole and fluorouracil. the north central cancer treatment group and the mayo clinic. *Journal of Clinical Oncology* **7,** 1447–1456.

Lavori, P. W. and Dawson, R. (2000). A design for testing clinical strategies: biased adaptive within-subject randomization. *Journal of the Royal Statistical Society: Series A (Statistics in Society)* **163,** 29–38.

Li, K.-C. (1991). Sliced inverse regression for dimension reduction. *Journal of the American Statistical Association* **86,** 316–327.

Lin, D. Y., Wei, L.-J., and Ying, Z. (1993). Checking the cox model with cumulative sums of martingale-based residuals. *Biometrika* **80,** 557–572.

Liu, X. and Zeng, D. (2013). Variable selection in semiparametric transformation models for right-censored data. *Biometrika* **100,** 859–876.

Lu, X., Chen, G., Singh, R. S., and K. Song, P. X. (2006). A class of partially linear single-index survival models. *Canadian Journal of Statistics* **34,** 97–112.

Lu, X., Singh, R., and Desmond, A. (2001). A kernel smoothed semiparametric survival model. *Journal of Statistical Planning and Inference* **98,** 119–135.

Luedtke, A. R. and van der Laan, M. J. (2016). Super-learning of an optimal dynamic treatment rule. *The International Journal of Biostatistics* **12,** 305–332.

Luo, S. and Ghosal, S. (2016). Forward selection and estimation in high dimensional single index models. *Statistical Methodology* **33,** 172–179.

Matsouaka, R. A., Li, J., and Cai, T. (2014). Evaluating marker-guided treatment selection strategies. *Biometrics* **70,** 489–499.

Moertel, C. G., Fleming, T. R., Macdonald, J. S., Haller, D. G., Laurie, J. A., Goodman, P. J., Ungerleider, J. S., Emerson, W. A., Tormey, D. C., Glick, J. H., et al. (1990). Levamisole and fluorouracil for adjuvant therapy of resected colon carcinoma. *New England Journal of Medicine* **322,** 352–358.

Murphy, S. A. (2003). Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **65,** 331–355.

Murphy, S. A. (2005). An experimental design for the development of adaptive treatment strategies. *Statistics in Medicine* **24,** 1455–1481.

Neumann, F.-J., Kastrati, A., Miethke, T., Pogatsa-Murray, G., Mehilli, J., Valina, C., Jogethaei, N., da Costa, C., Wagner, H., and Schomig, A. (2001). Treatment of chlamydia pneumoniae infection with roxithromycin and effect on neointima proliferation after coronary stent placement (ISAR-3): a randomised, double-blind, placebo-controlled trial. *Lancet* **357,** 2085–2089.

Nezhad, M. Z., Zhu, D., Li, X., Yang, K., and Levy, P. (2016). Safs: A deep feature selection approach for precision medicine. In *2016 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pages 501–506. IEEE.

Nielsen, J. P., Linton, O. B., et al. (1995). Kernel estimation in a nonparametric marker dependent hazard model. *The Annals of Statistics* **23,** 1735–1748.

O'Connell, M. J., Laurie, J. A., Kahn, M., Fitzgibbons Jr, R. J., Erlichman, C., Shepherd, L., Moertel, C. G., Kocha, W. I., Pazdur, R., Wieand, H. S., et al. (1998). Prospectively randomized trial of postoperative adjuvant chemotherapy in patients with high-risk colon cancer. *Journal of clinical oncology* **16,** 295–300.

Orellana, L., Rotnitzky, A., and Robins, J. M. (2010). Dynamic regime marginal structural mean models for estimation of optimal dynamic treatment regimes, part i: main content. *The International Journal of Biostatistics* **6,**.

Peng, H. and Huang, T. (2011). Penalized least squares for single index models. *Journal of Statistical Planning and Inference* **141,** 1362–1379.

Powell, J. L., Stock, J. H., and Stoker, T. M. (1989). Semiparametric estimation of index coefficients. *Econometrica: Journal of the Econometric Society* pages 1403–1430.

Radchenko, P. (2015). High dimensional single index models. *Journal of Multivariate Analysis* **139,** 266–282.

Robins, J., Orellana, L., and Rotnitzky, A. (2008). Estimation and extrapolation of optimal treatment and testing strategies. *Statistics in Medicine* **27,** 4678–4721.

Robins, J. M. (2004). Optimal structural nested models for optimal sequential decisions. In *Proceedings of the Second Seattle Symposium in Biostatistics*, pages 189–326. Springer.

Robins, J. M., Hernan, M. A., and Brumback, B. (2000). Marginal structural models and causal inference in epidemiology.

Robins, J. M., Rotnitzky, A., and Zhao, L. P. (1994). Estimation of regression coefficients when some regressors are not always observed. *Journal of the American Statistical Association* **89,** 846–866.

Royston, P. and Sauerbrei, W. (2008). Interactions between treatment and continuous covariates: a step toward individualizing therapy.

Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of educational Psychology* **66,** 688.

Sasieni, P. (1992a). Information bounds for the conditional hazard ratio in a nested family of regression models. *Journal of the Royal Statistical Society: Series B (Methodological)* **54,** 617–635.

Sasieni, P. (1992b). Non-orthogonal projections and their application to calculating the information in a partly linear cox model. *Scandinavian Journal of Statistics* pages 215–233.

Schumaker, L. (1981). *Spline Functions: Basic Theory.* Cambridge University Press.

Schwarz, G. et al. (1978). Estimating the dimension of a model. *The annals of statistics* **6,** 461–464.

Silverman, B. (1986). *Density Estimation for Statistics and Data Analysis*. Chapman & Hall/CRC Monographs on Statistics & Applied Probability. Taylor & Francis.

Song, R., Luo, S., Zeng, D., Zhang, H. H., Lu, W., and Li, Z. (2017). Semiparametric single-index model for estimating optimal individualized treatment strategy. *Electronic Journal of Statistics* **11,** 364.

Song, R., Wang, W., Zeng, D., and Kosorok, M. R. (2015). Penalized q-learning for dynamic treatment regimens. *Statistica Sinica* **25,** 901.

Steyerberg, E. W., Eijkemans, M. J., and Habbema, J. D. F. (1999). Stepwise selection in small data sets: a simulation study of bias in logistic regression analysis. *Journal of clinical epidemiology* **52,** 935–942.

Stoker, T. M. (1986). Consistent estimation of scaled coefficients. *Econometrica: Journal of the Econometric Society* pages 1461–1481.

Tibshirani, R. and Hastie, T. (1987). Local likelihood estimation. *Journal of the American Statistical Association* **82,** 559–567.

Tsiatis, A. (2007). *Semiparametric Theory and Missing Data*. Springer New York.

Tsiatis, A. A. (1981). A large sample study of cox's regression model. *The Annals of Statistics* pages 93–108.

Valsecchi, M., Silvestri, D., and Sasieni, P. (1996). Evaluation of long-term survival: use of diagnostics and robust estimators with Cox's proportional hazards model. *Statistics in medicine* **15,** 2763–2780.

van de Geer, S. (2000). *Empirical Processes in M-Estimation*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press.

van der Laan, M. J. and Luedtke, A. R. (2015). Targeted learning of the mean outcome under an optimal dynamic treatment rule. *Journal of Causal Inference* **3,** 61–95.

Van der Laan, M. J. and Rose, S. (2011). *Targeted learning: causal inference for observational and experimental data*. Springer Science & Business Media.

van der Vaart, A. W. and Wellner, J. A. (1996). *Weak Convergence and Empirical Processes*. Springer New York.

Vapnik, V. N. and Chervonenkis, A. Y. (1971). On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability & Its Applications* **16,** 264–280.

Wager, S. and Athey, S. (2018). Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association* **113,** 1228–1242.

Wang, W. (2004). Proportional hazards regression models with unknown link function and time-dependent covariates. *Statistica Sinica* pages 885–905.

Wasserman, L. (2006). *All of Nonparametric Statistics*. Springer New York.

Watkins, C. J. and Dayan, P. (1992). Q-learning. *Machine Learning* **8,** 279–292.

Watkins, C. J. C. H. (1989). *Learning from delayed rewards*. PhD thesis, King's College, Cambridge.

Weisberg, S. and Welsh, A. (1994). Adapting for the missing link. *The Annals of Statistics* pages 1674–1700.

Wolpert, D. H. (2002). *The supervised learning no-free-lunch theorems*. Springer.

Zemyan, S. M. (2012). *The Classical Theory of Integral Equations*. Birkhauser Boston.

Zeng, D. and Lin, D. Y. (2007). Maximum likelihood estimation in semiparametric regression models with censored data. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **69,** 507–564.

Zeng, D. and Lin, D. Y. (2014). Efficient estimation of semiparametric transformation models for two-phase cohort studies. *Journal of the American Statistical Association* **109,** 371–383.

Zhang, B., Tsiatis, A. A., Davidian, M., Zhang, M., and Laber, E. (2012). Estimating optimal treatment regimes from a classification perspective. *Stat* **1,** 103–114.

Zhang, B., Tsiatis, A. A., Laber, E. B., and Davidian, M. (2012). A robust method for estimating optimal treatment regimes. *Biometrics* **68,** 1010–1018.

Zhang, B., Tsiatis, A. A., Laber, E. B., and Davidian, M. (2013). Robust estimation of optimal dynamic treatment regimes for sequential treatment decisions. *Biometrika* **100,** 681–694.

Zhang, H. H. and Lu, W. (2007). Adaptive lasso for cox's proportional hazards model. *Biometrika* **94,** 691–703.

Zhao, Y., Kosorok, M. R., and Zeng, D. (2009). Reinforcement learning design for cancer clinical trials. *Statistics in Medicine* **28,** 3294–3315.

Zhao, Y., Zeng, D., Rush, A. J., and Kosorok, M. R. (2012). Estimating individualized treatment rules using outcome weighted learning. *Journal of the American Statistical Association* **107,** 1106–1118.

Zhao, Y., Zeng, D., Socinski, M. A., and Kosorok, M. R. (2011). Reinforcement learning strategies for clinical trials in nonsmall cell lung cancer. *Biometrics* **67,** 1422–1433.

Zhao, Y. Q., Zeng, D., Laber, E. B., Song, R., Yuan, M., and Kosorok, M. R. (2015). Doubly robust learning for estimating individualized treatment with censored data. *Biometrika* **102,** 151–168.

Zou, H. (2006). The adaptive lasso and its oracle properties. *Journal of the American statistical association* **101,** 1418–1429.