

# Variable Selection for Nonlinear Cox Regression Model via Deep Learning

Kexuan Li

Correspondence: Global Analytics and Data Sciences, Biogen, Cambridge, Massachusetts, USA

Received: November 16, 2022 Accepted: December 20, 2022 Online Published: December 24, 2022

doi:10.5539/ijsp.v12n1p21 URL: <https://doi.org/10.5539/ijsp.v12n1p21>

## Abstract

Variable selection problem for the nonlinear Cox regression model is considered. In survival analysis, one main objective is to identify the covariates that are associated with the risk of experiencing the event of interest. The Cox proportional hazard model is being used extensively in survival analysis in studying the relationship between survival times and covariates, where the model assumes that the covariate has a log-linear effect on the hazard function. However, this linearity assumption may not be satisfied in practice. In order to extract a representative subset of features, various variable selection approaches have been proposed for survival data under the linear Cox model. However, there exists little literature on variable selection for the nonlinear Cox model. To break this gap, we extend the recently developed deep learning-based variable selection model LassoNet to survival data. Simulations are provided to demonstrate the validity and effectiveness of the proposed method. Finally, we apply the proposed methodology to analyze a real data set on diffuse large B-cell lymphoma.

**Keywords:** variable selection, nonlinear Cox proportional hazards model, deep neural networks, survival analysis, LassoNet

## 1. Introduction

One of the areas of great methodologic advance in clinical trials and medical research has been the ability to handle censored time-to-event data. By censored time-to-event data, we mean that the value of observation is only partly known and the exact time-to-event is not observed. For example, a patient may drop off in the follow-up of the clinical study. For such censored data, we only know that the true time to event is less or greater than, or within the observed time, which refers to left censored, right censored and interval-censored respectively, and thus the classical statistical analysis methods may not be applicable to censored data. Survival analysis is a commonly-used method for the analysis of censored data, where the response is often referred to as a survival time, failure time, or event time. In practice, covariates are collected along with each observation object and the main objective of survival analysis is to demonstrate the dependence of the survival time on these covariates. To achieve this goal, the author in Cox (1972) introduced a famous semiparametric regression model that calculates the effect of each covariate on the survival time, which is called the Cox proportional hazards model. Instead of modeling the mean function of the survival time, the essential novelty of Cox proportional hazards is to model the hazard function, where the hazard function, as a function of  $t$ , is defined as instantaneous potential per unit time for the event to occur, given that the individual has survived up to time  $t$ . To be more specific, the classical Cox proportional hazard model assumes the logarithm of hazard is a linear function of the covariates, which, however, can be complex and nonlinear in practice. In contrast to the linear Cox proportional hazards model, a considerable interest has also been paid on nonlinear survival models, relaxing the assumption of linearity of log-hazard function (Fan et al. (1997); Ishwaran et al. (2008); Katzman et al. (2018))

In practice, not all the covariates may contribute to the prediction of the survival outcomes, and researchers are interested in identifying a subset of significant covariates upon which the hazard function depends. For example, in genome-wide association studies with time-to-event outcomes, people are interested in identifying the diagnostic single nucleotide polymorphisms that are associated with the response. From a statistical point of view, this kind of problem is referred to as variable selection or feature selection. Under the context of classical regression analysis, there are numbers of variable selection techniques, which are roughly categorized into three classes: filter methods, wrapper methods, and embedded methods. In the filter methods, features are filtered based on certain statistical measures for their correlation with the response, such as Pearson's correlation, and Fisher score. Wrapper methods evaluate different subsets of covariates on model and the best subset of covariates is selected based on the model performance; some commonly used wrapper methods include forward/backward selection or recursive feature elimination. Finally, in embedded methods, feature

selection process is embedded in the learning phase such as the least absolute shrinkage and selection operator (LASSO) (Tibshirani (1996)), the smoothly clipped absolute deviation (SCAD) (Fan and Li (2001)), and the adaptive LASSO (ALASSO) (Zou (2006)). The advantage of the embedded method is that the model training and feature selection can be performed simultaneously. See Saeys et al. (2007) for more details on the three methods. Many of these classical variable selection methods have been extended from regression analysis to survival analysis such as best subset variable selection, stepwise deletion, and Bayesian variable selection (Faraggi and Simon (1998); Ibrahim et al. (1999)). Penalized approach, which optimizes an objective function with a penalty function, has also been extended to right-censored failure time data (Fan and Li (2002); Tibshirani (1997); Zhang and Lu (2007)). More recently, the authors in Fan et al. (2010) extended the iterative sure feature screening procedure to Cox proportional hazards model. The works in Du et al. (2021); Yi et al. (2020); Zhao et al. (2019) considered the variable selection problem for interval-censored data. However, these methods rely on strong parametric assumptions that are often violated in practice. To overcome this problem, and owing to the advancement of deep learning for handling nonlinearity, we investigate the application of deep neural networks to variable selection for the nonlinear Cox model. To be more specific, we extend LassoNet proposed by Lemhadri et al. (2021) to right-censored survival data.

Besides the above three categories of variable selection methods, in recent years, applying deep learning to variable selection has also made a great breakthrough. To name a few, the work of Han et al. (2018); Liu et al. (2017) added row-sparse regularization to hidden layers to achieve feature selection; the authors in Abid et al. (2019) applied reparametrization trick, or Gumbel-softmax trick to autoencoder and obtain discrete feature selection; in Mirzaei et al. (2020), the authors established a teacher-student neural network for feature selection; LassoNet, proposed by Lemhadri et al. (2021), introduced a novel feature selection framework by penalizing the parameters in the residual layer with the constraint that the norm of the parameters in the first layer is less than the corresponding norm of the parameters in the residual layer; the method introduced in Li et al. (2022) combined deep neural networks and feature screening under the high-dimensional, low-sample-size setting. Even though these deep learning-based variable selection methods have made great success in regression or classification problems, the extension of them to survival data is not trivial. One of the objectives of this paper is to break this gap.

The rest of the paper is organized as follows. In Section 2, we formulate the problem of interest and briefly review the recent applications of deep learning in survival analysis. Section 3 describes the variable selection procedure in the nonlinear Cox regression model using LassoNet. Section 4 gives several comprehensive simulation studies to evaluate finite sample performances of the proposed methods. Finally, in Section 5, we apply our method to diffuse large B-cell lymphoma and give our conclusion in Section 6.

## 2. Problem Formulation and Related Works

In this section, we first formulate the problem of interest and then provide a literature review of recent deep learning models applied in survival analysis.

### 2.1 Problem Formulation

Denote  $T, C, \mathbf{X} \in \mathbb{R}^p$  by the survival time, the censoring time, and the corresponding covariates, respectively. For simplicity, we assume that  $T$  and  $C$  are conditionally independent given  $\mathbf{X}$ . In practice,  $T$  may not be completely observed, and let  $Y = \min\{T, C\}$ ,  $\delta = \mathbb{1}(T \leq C)$  be the observed time and the censoring indicator, respectively. Thus, our observed data set is  $\{(Y_i, \delta_i, \mathbf{x}_i) : Y_i \geq 0, \delta_i \in \{0, 1\}, \mathbf{x}_i \in \mathbb{R}^p, i = 1, \dots, n\}$  which is independently and identically distributed from  $(Y, \delta, \mathbf{X})$ . Furthermore, we let  $\mathcal{R}(t) = \{i : Y_i \geq t\}$  be the set of all individuals who are available (have survived) at time  $t$ , which is referred to as the risk set. In survival analysis, with a particular value  $\mathbf{x}$  for the covariate  $\mathbf{X}$ , there are two commonly used quantitative measurements to describe the distribution of survival time  $T$ . The first one is the survival function  $S(t|\mathbf{x})$ , which is defined as  $S(t|\mathbf{x}) = \mathbb{P}(T > t|\mathbf{X} = \mathbf{x})$ , representing the probability of surviving beyond time  $t$ . Another one is the hazard function  $h(t|\mathbf{x})$ , which is defined as:

$$h(t|\mathbf{x}) = \lim_{\Delta t \rightarrow 0} \frac{\mathbb{P}(t \leq T < T + \Delta t | T \geq t, \mathbf{X} = \mathbf{x})}{\Delta t}.$$

According to the definition of  $h(t|\mathbf{x})$ , the hazard function characterizes the instantaneous rate of failure at time  $t$  given that the individual has survived up to time  $t$ . Given an observed data set  $\{(Y_i, \delta_i, \mathbf{x}_i)\}$ , the corresponding complete likelihood function of the observed data is given by

$$L_{\text{complete}} = \prod_{i \in \mathcal{U}} f(Y_i|\mathbf{x}_i) \prod_{i \in \mathcal{C}} S(Y_i|\mathbf{x}_i) = \prod_{i \in \mathcal{U}} h(Y_i|\mathbf{x}_i) \prod_{i=1}^n S(Y_i|\mathbf{x}_i), \tag{1}$$

where  $\mathcal{U} = \{i : \delta_i = 1\}$  and  $\mathcal{C} = \{i : \delta_i = 0\}$  are the index set of uncensored and censored data respectively, and  $f(t|\mathbf{x})$  is the conditional density function. In practice, people are interested in modeling the possible relationship between the survival

time  $T_i$  and the covariates  $\mathbf{x}_i$ . The Cox proportional hazards model in Cox (1972) is a popular approach to achieve this goal. To be more specific, given features  $\mathbf{x} = (x_1, \dots, x_p)^\top$ , Cox proportional hazards model assumes that the covariate has a log-linear effect on the hazard function, that is,

$$h(t|\mathbf{x}) = h_0(t) \exp[\mathbf{x}^\top \boldsymbol{\beta}] = h_0(t) \exp\left[\sum_{j=1}^p x_j \beta_j\right], \tag{2}$$

where  $h_0(t)$  is a completely unspecified baseline hazard function assuming to be constant across all individuals and  $\boldsymbol{\beta} = (\beta_1, \dots, \beta_p)^\top$  is the regression coefficient. The model assumes that the hazard function  $h(t|\mathbf{x})$  is a multiplication of a non-parametric hazard function and a relative risk function, which is given by  $\mathbf{x}^\top \boldsymbol{\beta}$ . Let  $H_0(t) = \int_0^t h_0(t)$  be the cumulative baseline hazard function. Under Cox proportional hazards model in (2), using the fact that  $S(t|\mathbf{x}) = \exp[-\int_0^t h(t|\mathbf{x})]$ , the complete likelihood function in (1) becomes

$$L_{\text{complete}}(\boldsymbol{\beta}) = \prod_{i \in \mathcal{U}} h_0(Y_i) \exp[\mathbf{x}_i^\top \boldsymbol{\beta}] \prod_i^n \exp\{-H_0(Y_i) \exp[\mathbf{x}_i^\top \boldsymbol{\beta}]\}. \tag{3}$$

Specifically, if  $h_0(t)$  (or equivalent  $H_0(t)$ ) is parameterized by  $\theta$ , then both  $\theta$  and  $\boldsymbol{\beta}$  can be estimated by maximizing the complete likelihood function in (3). However, in practice, the form of  $h_0(t)$  is unknown, thus, instead of optimizing  $L_{\text{complete}}(\boldsymbol{\beta})$ ,  $\boldsymbol{\beta}$  can be estimated by maximizing the following partial likelihood function, without specification of  $h_0(t)$

$$L_{\text{partial}}(\boldsymbol{\beta}) = \prod_{i=1}^n \left( \frac{\exp[\mathbf{x}_i^\top \boldsymbol{\beta}]}{\sum_{j \in \mathcal{R}(Y_i)} \exp[\mathbf{x}_j^\top \boldsymbol{\beta}]} \right)^{\delta_i}. \tag{4}$$

From now on, we let  $L(\boldsymbol{\beta}) = L_{\text{partial}}(\boldsymbol{\beta})$  and  $\ell(\boldsymbol{\beta}) = \log L(\boldsymbol{\beta})$  for simplicity. Notice that the model in (2) assumes the covariate effect on the hazard function is log-linear, however, the linearity assumption is too stringent in real-world applications and thus we should consider more flexible choices for the covariate effect to model the nonlinearity and interaction.

Now, we describe the problem of variable selection in right-censored data. In variable selection, the goal is to identify a subset  $\mathcal{S} \subseteq \{1, 2, \dots, p\}$  of the most discriminative and informative features with size  $|\mathcal{S}| = k \leq p$  and also a function  $g : [0, \infty) \times \mathbb{R}^k \rightarrow \mathbb{R}$  such that the hazard function  $h(t|\mathbf{x})$  in (2) equals  $g(t|\mathbf{x}^{(\mathcal{S})})$ , where  $\mathbf{x}^{(\mathcal{S})} = (x_{s_1}, \dots, x_{s_k})^\top \in \mathbb{R}^k$  and  $s_i \in \mathcal{S}$ .

### 2.2 Deep Learning for Survival Analysis

Recently, deep learning has made great breakthroughs in a wide range of applications, such as natural language processing (Bahdanau et al. (2014)), computer vision (He et al. (2016)), dynamics system (Li et al. (2021)), drug discovery and toxicology (Jiménez-Luna et al. (2020)). Owing to the superior performance and good theoretical guarantees of deep learning, applying deep learning to survival data has also drawn much attention. In this subsection, we provide a brief review of recent contributions of deep learning in the field of survival analysis.

To the best of our knowledge, the first attempt to apply deep neural networks in survival data is by Faraggi and Simon (1995), where the authors extended the linear Cox proportional hazards (Cox (1972)) by replacing the linear combination of covariates in (2) with a single layer feed-forward neural network. However, they showed that the proposed model fail to outperform the classical linear Cox model. Under a similar framework, instead of using a single-layer neural network, Katzman et al. (2018) explored more complex networks and demonstrated that modern deep learning techniques perform as well as or better than the linear Cox model. In addition,  $\ell_1$  and  $\ell_2$  regularization terms have been utilized to the loss function to reduce the over-fitting. Kvamme et al. (2019) further alleviated the proportional hazards assumption by approximating the loss function based on case-control sampling. The authors also came up with a heuristic to approximate the gradients, enabling the models to scale to large data sets. Ren et al. (2019) considered deep recurrent neural network to capture the sequential patterns of the feature over time in survival analysis. Lee et al. (2018) introduced a novel approach called DeepHit which is able to directly estimate the joint distribution of events in discrete time. To improve discriminative performance, Lee et al. (2018) combined the log-likelihood function with a ranking loss and is applicable for competing risks. Zhu et al. (2017) added convolutional architectures to the loss function in order to capture image-based covariates and applied the method to pathological images of lung cancer. Luck et al. (2017) combined Cox partial likelihood function with an isotonic regression loss to improve the predictive performance. Manduchi et al. (2021) studied the problem of clustering survival data by using a deep generative model. Nagpal et al. (2021a) introduced a fully parametric approach, named deep survival machines, to estimate relative risks in time-to-event prediction problems by modeling the survival function as a weighted mixture of individual parametric survival distributions. Nagpal et al. (2021b) further extended their

work to a mixture of Cox model, by assuming there exists a latent group structure and within each group, the proportional hazards assumption holds. Zhong et al. (2021) considered deep neural networks to generalize the extended hazard model, which includes both the Cox proportional hazards model and the accelerated failure time model as special cases, and theoretically proved the consistency as well as the rate of convergence of the proposed estimator. However, all the works mentioned above failed to touch the variable selection problem in survival, and the main objective of this paper is to fill the gap.

### 3. Methods

Before introducing our method, let us briefly review the classical penalized partial likelihood techniques for Cox proportional hazards model. To the best of our knowledge, Tibshirani (1996) is the first attempt to use the penalization method to achieve variable selection on survival data, where the author adds an  $\ell_1$  penalty term in the partial likelihood function to estimate  $\beta$ , that is,

$$\hat{\beta}_{\text{lasso}} = \operatorname{argmin} -\ell(\beta), \text{ subject to } \sum_{j=1}^p |\beta_j| \leq s, \tag{5}$$

where  $s$  is a tuning parameter. The above equation can be rewritten equivalently in a Lagrangian form

$$\hat{\beta}_{\text{lasso}} = \operatorname{argmin} -\ell(\beta) + \lambda \sum_{j=1}^p J(\beta_j), \tag{6}$$

where  $J(\beta_j) = |\beta_j|$  and the exact relationship between  $s$  and  $\lambda$  is data dependent. In addition to the  $\ell_1$  penalty, some other penalty forms have also been studied in the literature, such as adaptive Lasso (Zhang and Lu (2007)), SCAD (Fan and Li (2002)). However, the aforementioned variable selection methods ignore feature interaction and nonlinear structure, to overcome the problem, we extend the advantages of LassoNet (Lemhadri et al. (2021)) to survival data to tackle nonlinearity and complex structure. Before presenting the proposed variable selection approach using LassoNet, we first briefly review the notations and terminologies used in deep neural networks. Unlike traditional approximation theory in mathematics where people approximate complex functions based upon the summation of a series of simple functions, such as the Taylor Series and Fourier series, deep learning uses the compositions of simple functions. To be more specific, from a high-level point of view, a typical feed-forward neural network  $g(x, \mathbf{W}, \mathbf{v})$  is a composition of several simple functions, that is,

$$g(x, \mathbf{W}, \mathbf{v}) := \mathbf{W}_L \sigma_{v_L} \dots \mathbf{W}_1 \sigma_{v_1} \mathbf{W}_0 x, \quad x \in \mathbb{R}^{d_0}, d_0 = p, \tag{7}$$

where  $\mathbf{W} \in \mathcal{W} := \{(\mathbf{W}_0, \dots, \mathbf{W}_L) : \mathbf{W}_l \in \mathbb{R}^{d_{l+1} \times d_l}, 0 \leq l \leq L\}$  is the weight matrix,  $\mathbf{v} \in \mathcal{V} := \{(v_1, \dots, v_L) : v_l \in \mathbb{R}^{d_l}, 1 \leq l \leq L\}$  is the bias term, and  $\sigma_v : \mathbb{R}^r \rightarrow \mathbb{R}^r$  is a non-linear activation function to learn the complex pattern from the data. Here  $L$  is the number of hidden layers, i.e., the length of the network, and  $\mathbf{d} = (d_0 = p, \dots, d_{L+1}), d_{L+1} = 1, d_j > 0$  is the number of units in each layer, i.e., the depth of the network. In our paper, we consider the following shift activation function

$$\sigma_v \begin{pmatrix} y_1 \\ \vdots \\ y_r \end{pmatrix} = \begin{pmatrix} \sigma(y_1 - v_1) \\ \vdots \\ \sigma(y_r - v_r) \end{pmatrix},$$

where  $\mathbf{v} = (v_1, \dots, v_r) \in \mathbb{R}^r$ . The parameters need to be estimated from the data are  $(\mathbf{W}_j)_{j=0, \dots, L}$  and  $(v_j)_{j=1, \dots, L}$ . In practice, the network depth is of crucial importance to capture the complex hidden structure of the data. A deeper network is able to learn a more complex representation of the input data, however, with the network depth increasing, accuracy gets saturated and then degrades rapidly, this is called degradation problems. To overcome this problem, He et al. (2016) introduced a famous framework, denoted as residual neural networks. By adding skip connections, the number of layers in residual neural networks can easily achieve more than one hundred. Thus, we consider the following class of residual feed-forward neural networks

$$\mathcal{F} = \{f(x, \mathbf{W}, \mathbf{v}, \theta) := \theta^\top x + g(x, \mathbf{W}, \mathbf{v}), \quad g(x, \mathbf{W}, \mathbf{v}) \text{ of form (7)}, \quad \theta, x \in \mathbb{R}^p. \tag{8}$$

To tackle the complicated structure between the covariates and survival time, we replace the linear function in (2) by a complex nonlinear function  $\Phi(x, \mathbf{W}, \mathbf{v}, \theta) \in \mathcal{F}$  and the corresponding hazard function and partial likelihood function now becomes

$$h(t|x) = h_0(t) \exp[\Phi(x, \mathbf{W}, \mathbf{v}, \theta)], \tag{9}$$

$$L(\mathbf{W}, \mathbf{v}, \theta) = \prod_{i=1}^n \left( \frac{\exp[\Phi(x_i, \mathbf{W}, \mathbf{v}, \theta)]}{\sum_{j \in \mathcal{R}(Y_i)} \exp[\Phi(x_j, \mathbf{W}, \mathbf{v}, \theta)]} \right)^{\delta_i}. \tag{10}$$

By replacing the linear relationship  $\mathbf{x}^\top \boldsymbol{\beta}$  to  $\Phi(\mathbf{x}, \mathbf{W}, \mathbf{v}, \boldsymbol{\theta})$ , the approach is capable of modeling interactions and nonlinearity between the covariates and survival time. In order to achieve variable selection, penalty terms are needed in the following optimization problem

$$(\widehat{\mathbf{W}}, \widehat{\mathbf{v}}, \widehat{\boldsymbol{\theta}}) = \operatorname{argmin} L(\mathbf{W}, \mathbf{v}, \boldsymbol{\theta}) + \lambda \|\boldsymbol{\theta}\|_1, \tag{11}$$

$$\text{subject to } \|W_{0,i}\|_\infty \leq M|\theta_i|, i = 1, \dots, p. \tag{12}$$

In (11),  $\theta_i$  is the  $i$ -component of  $\boldsymbol{\theta} \in \mathbb{R}^p$  and  $W_{0,i}$  contains the weights for the  $i$ -th input variable in the first hidden layer. There are two tuning parameters in the objective function:  $\lambda$  and  $M$ , which penalize the linear and nonlinear components simultaneously. The constraint  $\|W_{0,i}\|_\infty \leq M|\theta_i|$  plays a vital role, in that  $M$  leverages the effect of the  $i$ -th input variable, thereby capturing the non-linearity in the data. When  $M = 0$ , the neural network part  $g(\mathbf{x}, \mathbf{W}, \mathbf{v})$  in  $\Phi(\mathbf{x}, \mathbf{W}, \mathbf{v}, \boldsymbol{\theta})$  vanishes and  $\Phi(\mathbf{x}, \mathbf{W}, \mathbf{v}, \boldsymbol{\theta}) = \boldsymbol{\theta}^\top \mathbf{x}$ , and thus (11) degenerates to standard LASSO in (Tibshirani (1996)). When  $M \rightarrow \infty$ , we get a feed-forward network with an  $\ell_1$  penalty on the residual layer. Following the optimization procedure outlined in Lemhadri et al. (2021), we optimize the objective function (11) using proximal gradient descent. It is worth mentioning that without domain knowledge, it is difficult to determine the value of  $M$ . In practical implementation, we use cross-validation to select  $M$ , a common practice in machine learning algorithms. By introducing the residual layer  $\boldsymbol{\theta}^\top \mathbf{x}$ , the linear and non-linear components can be optimized simultaneously. The architecture of the network and the detailed pseudocode are summarized in Figure 1 and Algorithm 1.

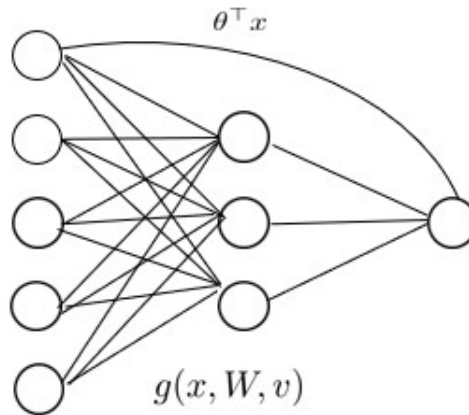


Figure 1. The architecture of LassoNet.  $\boldsymbol{\theta}^\top \mathbf{x}$  is a residual layer and  $g(\mathbf{x}, \mathbf{W}, \mathbf{v})$  is a feed-forward neural network

---

**Algorithm 1**

**Input:** observed data set is  $\{(Y_i, \delta_i, \mathbf{x}_i) : Y_i \geq 0, \delta_i \in \{0, 1\}, \mathbf{x}_i \in \mathbb{R}^p, i = 1, \dots, n\}$ , residual network  $\Phi(\mathbf{x}, \mathbf{W}, \mathbf{v}, \boldsymbol{\theta}) \in \mathcal{F}$ , number of epochs  $E$ , learning rate  $\alpha$ , hyper-parameter  $M$ , and path multiplier  $\epsilon$

Initialize  $\boldsymbol{\theta}, \mathbf{W}, \mathbf{v}, \lambda, k = p$ , and loss function  $L(\mathbf{W}, \mathbf{v}, \boldsymbol{\theta}) = \prod_{i=1}^n \left( \frac{\exp[\Phi(\mathbf{x}, \mathbf{W}, \mathbf{v}, \boldsymbol{\theta})]}{\sum_{j \in \mathcal{R}(Y_i)} \exp[\Phi(\mathbf{x}, \mathbf{W}, \mathbf{v}, \boldsymbol{\theta})]} \right)^{\delta_i}$

**while**  $k > 0$  **do**

    Update  $\lambda \leftarrow (1 + \epsilon)\lambda$

**for**  $e \in \{1, \dots, E\}$  **do**

        Compute gradients  $\nabla_{\boldsymbol{\theta}} \ell, \nabla_{\mathbf{W}} \ell, \nabla_{\mathbf{v}} \ell$  using back-propagation

        Update  $\boldsymbol{\theta} \leftarrow \boldsymbol{\theta} - \alpha \nabla_{\boldsymbol{\theta}} \ell, \mathbf{W} \leftarrow \mathbf{W} - \alpha \nabla_{\mathbf{W}} \ell, \mathbf{v} \leftarrow \mathbf{v} - \alpha \nabla_{\mathbf{v}} \ell$

        Update  $(\boldsymbol{\theta}, W_0) \leftarrow \text{Hier-Prox}(\boldsymbol{\theta}, W_0, \alpha \lambda, M)$ , where

        Hier-Prox is provided in the Algorithm 2 in Lemhadri et al. (2021)

**end for**

    Update  $k$  to be the number of non-zero elements of  $\boldsymbol{\theta}$

**end while**

---

#### 4. Numerical Results

In this section, we carry out simulation studies to assess the finite sample performance of the LassoNet on right-censored survival data. In both designs, we use the same network architecture with length  $L = 3$  and width  $N = 30$ . The dropout probability is set to 0.2 to avoid overfitting. Summary statistics from each simulation setting are calculated based on 100 independent simulation runs. In each setting, we use the inverse probability method by Bender et al. (2005) to generate the event times from the hazard function. Suppose the conditional survival function under Cox proportional hazards model in (2) is

$$S(t|\mathbf{x}) = \exp(-H_0(t) \exp(\psi(\mathbf{x}))),$$

where  $H_0(t)$  is the cumulative baseline hazard function and  $\psi(\mathbf{x})$  is the covariates effect. let  $U$  be uniformly distributed on  $[0, 1]$ , then the corresponding event time

$$T = S^{-1}(U|\mathbf{x}) = H_0(t)^{-1} \left( -\frac{\log(U)}{\exp(\psi(\mathbf{x}))} \right). \tag{13}$$

In order to evaluate the performance of the variable selection method, we consider the following three metrics:

- MinSize = the minimum number of selected variables to includes all true variables
- Prob( $k$ , all) = the success rate that the selected top  $k$  variables are all true variables
- Prob( $k$ ,  $i$ ) = the success rate that the selected top  $k$  variables contain variable  $x_i$

These three metrics are widely used in feature selection literature. Suppose the number of true variables is  $k$ , then by definition, MinSize is expected to be at least  $k$ ; and the closer to  $k$ , the better the procedure. The metric Prob( $k$ , all) measures the sensitivity (true positive) rate of the method in detecting the true variables. A higher value of Prob( $k$ , all) is desirable, as this indicates there is a higher chance that the algorithm will pick all the true variables. The metric Prob( $k$ ,  $i$ ) measures the true positive and false positive rate of the proposed method. If the set of the true variable contains variable  $x_i$ , then the higher value of Prob( $k$ ,  $i$ ), the better of the model, in contrast, if  $x_i$  is not associated with the survival time, then the lower value of Prob( $k$ ,  $i$ ), the better of the model.

We compare the proposed model  $\text{COX}_{\text{LassoNet}}$  with three other methods: (1) classical Cox proportional hazards model  $\text{COX}_{\text{classical}}$ ; (2)  $\ell_1$  regularized Cox model  $\text{COX}_{\text{Lasso}}$ ; and (3) stepwise variable selection Cox proportional hazards model  $\text{COX}_{\text{Stepwise}}$ . In each method, we rank the importance of each feature and select the top  $k$  features accordingly. The importance of each feature from model  $\text{COX}_{\text{classical}}$  is based on the  $p$ -value associated with each variable. The  $\ell_1$  regularized Cox model  $\text{COX}_{\text{Lasso}}$  and stepwise variable selection Cox proportional hazards model  $\text{COX}_{\text{Stepwise}}$  are implemented by the “**coxphMIC**” package and “**pec**” package in R.

##### 4.1 Model 1: Linear Cox Proportional Hazards Model

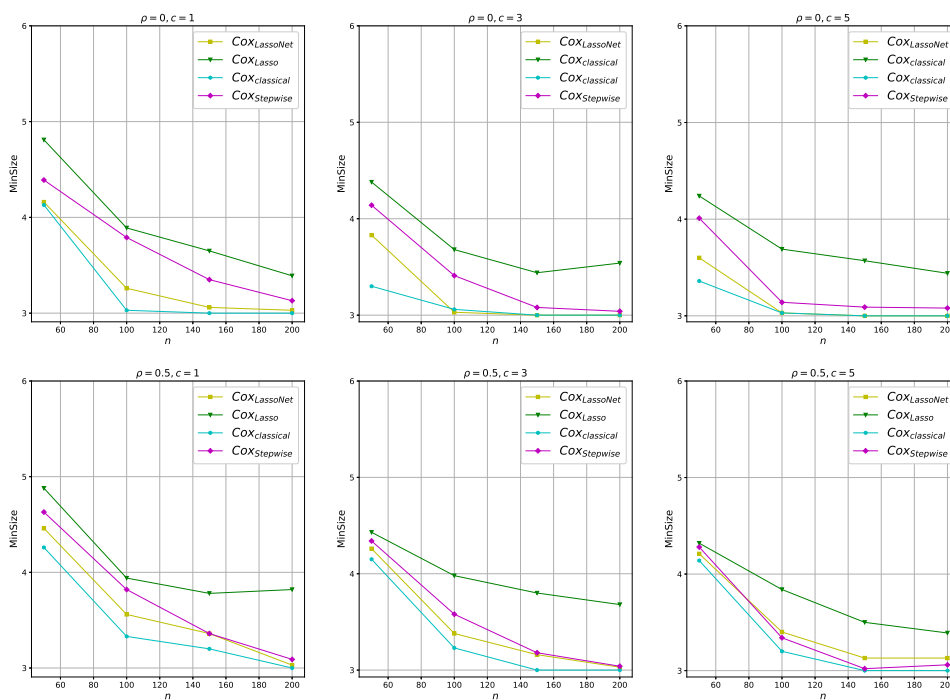
In the first model, we consider the classical Cox proportional hazards model in (2) with  $\boldsymbol{\beta} = (0.8, 0, 0, 1, 0, 0, 0, 0, 0.6, 0)^\top$ , and the ten covariates  $\mathbf{x} = (x_1, \dots, x_{10})$  are from multivariate Gaussian distribution with piecewise correlation  $\text{Cor}(x_i, x_j) = \rho^{|i-j|}$ ,  $i, j = 1, \dots, 10$ . Event times are generated from (13) with  $\psi(\mathbf{x}) = \mathbf{x}^\top \boldsymbol{\beta}$  and without loss of generality, we let  $h_0(t) = 1$ . Censoring times  $C$  are generated from a uniform distribution  $U(0, c)$ , where  $c$  is a predetermined parameter to control the rates of censored samples. The MinSize and Prob(3, all) are summarized in Figure 2, which suggests that the MinSize decreases and Prob(3, all) increases as  $n$  increases. Moreover, as  $c$  increases, the results get better. This is because increasing  $c$  will diminish the impact of censoring. Furthermore, we investigate the influence of  $\rho$ , i.e., the correlation of the covariates. We observe that when the dependence increases ( $\rho$  increases), it is more difficult to select the true variables for all the methods. From Figure 2, we also find that classical Cox regression outperforms other methods. This is mainly because the data generating process is nothing but the linear Cox proportional hazards model thus the parametric test is more powerful than other nonparametric methods. Among all the other three methods  $\text{COX}_{\text{LassoNet}}$ ,  $\text{COX}_{\text{Lasso}}$ , and  $\text{COX}_{\text{Stepwise}}$ , the deep learning method is the best. The results of Prob( $k$ ,  $i$ ),  $i = 1, 4, 9$  is reported in Figure 3 to check the probability of selecting each true variable for each method. It is straightforward to see that,  $\text{Prob}(3, 4) \geq \text{Prob}(3, 1) \geq \text{Prob}(3, 9)$ , which is consistent with the fact that the coefficient of  $x_4$  is greater than  $x_1$  and  $x_9$ . We only report the results for  $\rho = 0$  to save space, and the pattern for  $\rho = 0.5$  is similar.

##### 4.2 Model 2: Nonlinear Cox Proportional Hazards Model

In the second example, we consider a nonlinear Cox proportional hazards where the hazard function is generated by

$$h(t) = x_1 + \max(x_4, 1) + x_4 \times x_9.$$

MinSize



Prob(3, all)

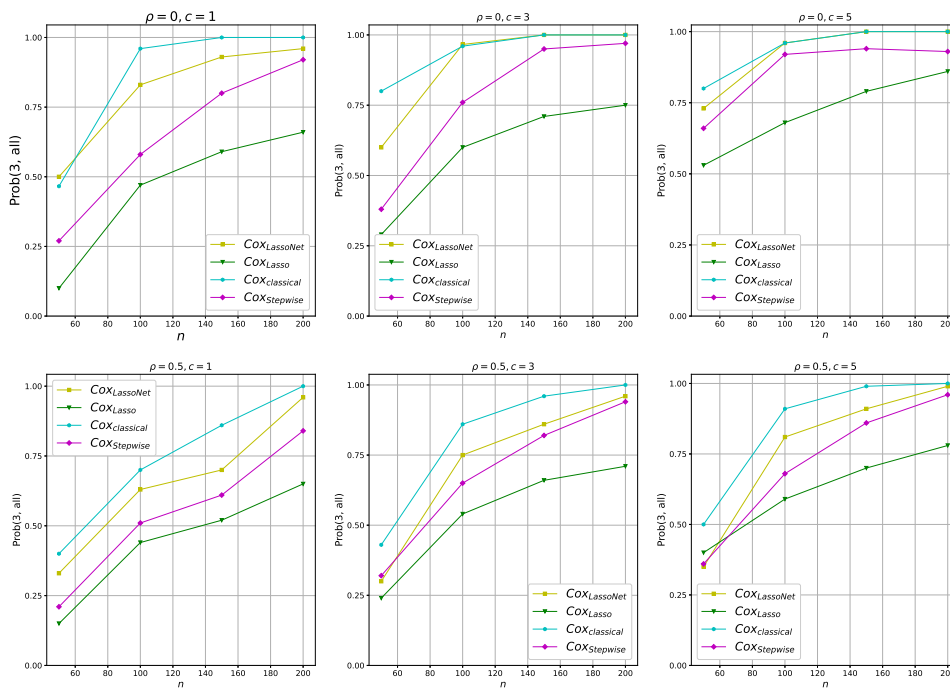


Figure 2. *MinSize and Prob(3, all) of Simulation 1. The true hazard function is a linear function defined as  $h(t) = 0.8x_1 + x_4 + 0.6x_9$ . LassoNet:  $COX_{LassoNet}$ ; Lasso:  $COX_{Lasso}$ ; Cox:  $COX_{classical}$ ; Piecewise:  $COX_{Stepwise}$*

Similar to Simulation 1, we generate the ten covariates  $\mathbf{x} = (x_1, \dots, x_{10})$  are from multivariate Gaussian distribution with piecewise correlation  $Cor(x_i, x_j) = \rho^{|i-j|}$ ,  $i, j = 1, \dots, 10$  and the baseline hazard function  $h_0(t) = 1$ . In this example, variable  $x_1$  is a linear term while  $x_4, x_9$  are nonlinear and interacted terms. The goal of this setting is to compare the ability of selecting the nonlinear variables. The summary statistics of different metrics are summarized in Figure 4 and

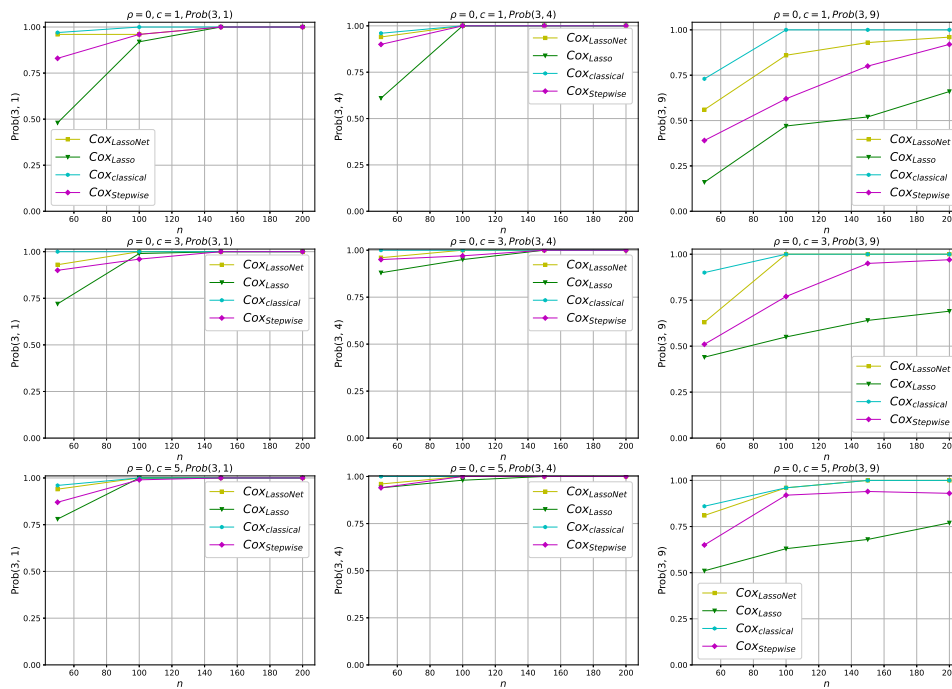


Figure 3.  $Prob(3, 1)$ ,  $Prob(3, 4)$ , and  $Prob(3, 9)$  of Simulation 1. The true hazard function is a linear function defined as  $h(t) = 0.8x_1 + x_4 + 0.6x_9$

Figure 5. It is straightforward to see that LassoNet is more capable of selecting the nonlinear features than the other three methods. It is worth mentioning that  $Prob(3, 1)$  is very closed to 1 for all the methods, however,  $Prob(3, 4)$  and  $Prob(3, 9)$  are very low for all the other methods.

**5. Real Data Analysis**

In this section, we apply the proposed methods to the diffuse large-B-cell lymphoma data (Rosenwald et al. (2002)). The data set consists of  $n = 240$  observations with diffuse large-B-cell lymphoma after chemotherapy and 138 patients die during the follow-up study, that is, the censor rate is  $138/240 = 57.5\%$ . DNA microarrays techniques were examined for gene expression to analyze for genomic abnormalities. The purpose of this analysis is to identify the prognostic genes among  $p = 7399$  microarray features. The length and width of the network are  $L = 3, N = 30$ , respectively with a dropout rate equal 0.2.

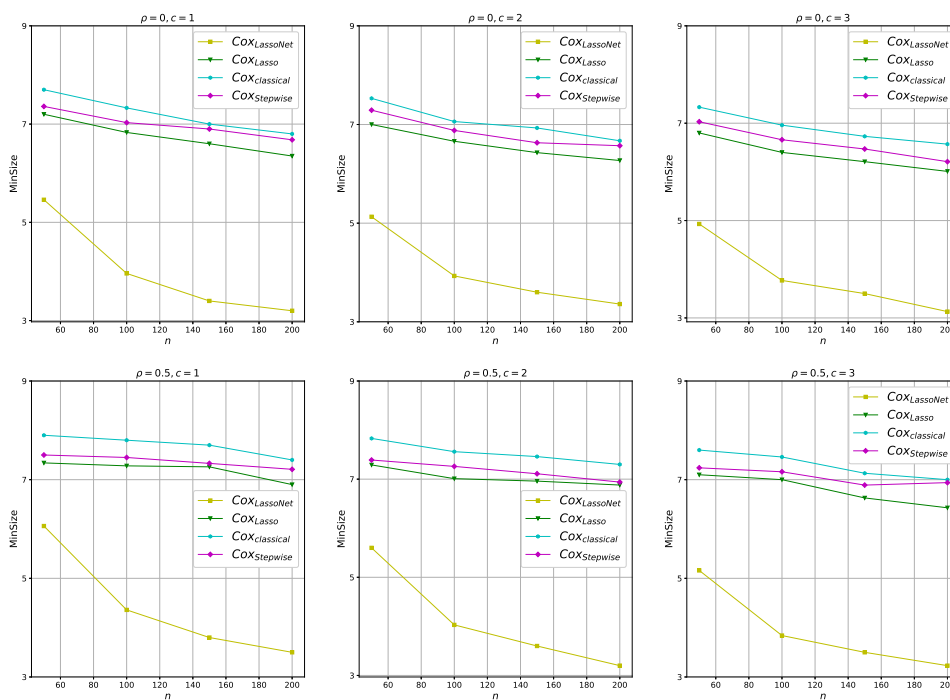
We standardize each predictor to have a mean zero and standard deviation 1 and apply LassoNet to select the top five genes, which are: BC012161, LC33732, X00457, AK000978, and LC24433. It is interesting that three of the selected genes, BC012161, X00457, and LC24433 belong to the proliferating cells, major-histocompatibility-complex class, and immune cells in the lymph node, respectively. These three classes were the gene-expression signatures defined in Rosenwald et al. (2002). The other two selected genes do not belong to the classes defined in Rosenwald et al. (2002), however, they may also be related to lymphoma as reported by Sha et al. (2006). To further verify the selected genes are significant, we fit the Cox proportional hazard model with all these five genes and report the  $p$ -values. The  $p$ -values are all less than  $10^{-4}$  which confirms that the five genes are significant. The results are summarized in Table 5.

Table 1. Top five genes selected by LassoNet for diffuse large-B-cell lymphoma data

Genes	Description	$p$ -value
BC012161	reported by Rosenwald et al. (2002)	$< 10^{-5}$
LC33732	reported by Sha et al. (2006)	$< 10^{-5}$
X00457	reported by Rosenwald et al. (2002)	$< 10^{-5}$
AK000978	reported by Sha et al. (2006)	$< 10^{-5}$
LC24433	reported by Rosenwald et al. (2002)	$< 10^{-4}$



MinSize



Prob(3, all)

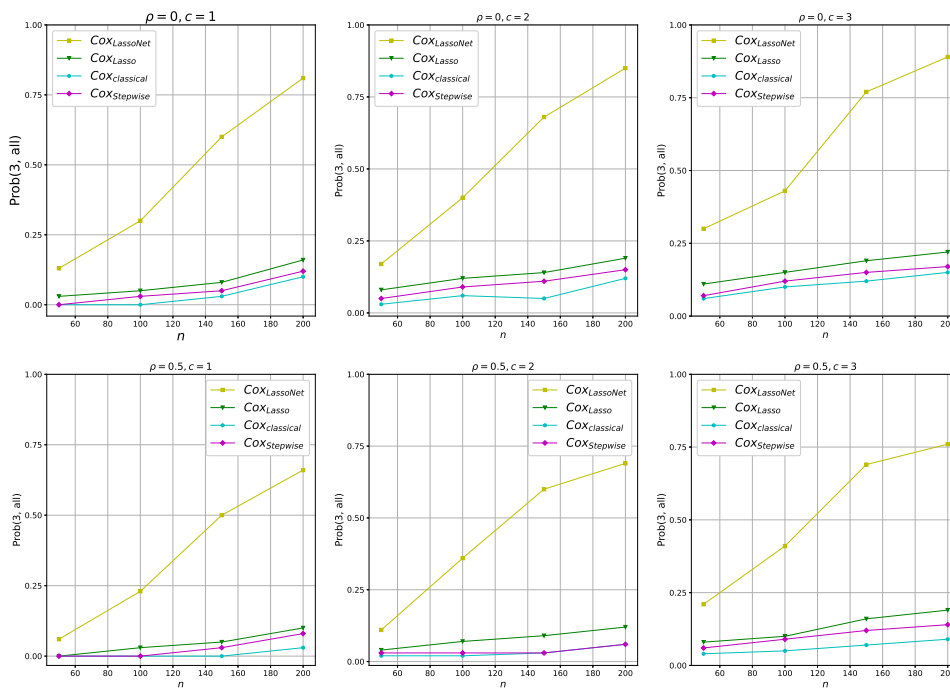


Figure 4. MinSize and Prob(3, all) of Simulation 2. The true hazard function is a nonlinear function defined as  $h(t) = x_1^2 + \max(x_4, 1) + x_4 \times x_9$

6. Conclusion

The successful application of deep neural networks in variable selection problem on survival data is limited, In this work, we extend the state-of-the-art deep learning variable selection algorithm to survival data. Owing to the superiority of deep neural networks to capture complex structure feature space, the proposed method is able to select the informative features

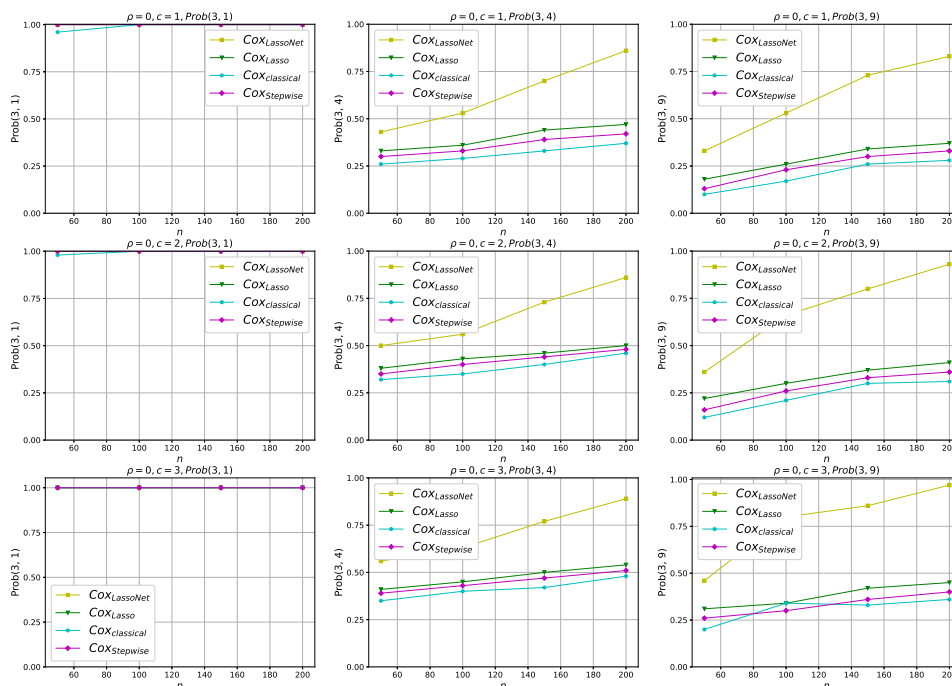


Figure 5.  $Prob(3, 1)$ ,  $Prob(3, 4)$ , and  $Prob(3, 9)$  of Simulation 2. The true hazard function is a nonlinear function defined as  $h(t) = x_1^2 + \max(x_4, 1) + x_4 \times x_9$

when the true hazard function is nonlinear. Simulation results verify the validity of the method and the real world data analysis indicates that the deep learning method can also be applied to survival data with the dimensionality that can be much larger than sample size.

**Acknowledgement**

The authors are grateful for Peng Sun and Richard Foster for their many critical and constructive comments and suggestions that greatly improved the paper.

**References**

Abid, A., Balin, M. F., & Zou, J. (2019). Concrete autoencoders for differentiable feature selection and reconstruction. *arXiv preprint arXiv:1901.09346*.

Bahdanau, D., Cho, K., & Bengio, Y. (2014). Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.

Bender, R., Augustin, T., & Blettner, M. (2005). Generating survival times to simulate cox proportional hazards models. *Statistics in medicine*, 24(11), 1713–1723.

Cox, D. R. (1972). Regression models and life-tables. *Journal of the Royal Statistical Society: Series B (Methodological)*, 34(2), 187–202.

Du, M., Zhao, H., & Sun, J. (2021). A unified approach to variable selection for coxs proportional hazards model with interval-censored failure time data. *Statistical Methods in Medical Research*, 30(8), 1833–1849.

Fan, J., Feng, Y., & Wu, Y. (2010). High-dimensional variable selection for coxs proportional hazards model. In *Borrowing strength: Theory powering applications—a Festschrift for Lawrence D. Brown*, pages 70–86. Institute of Mathematical Statistics.

Fan, J., Gijbels, I., & King, M. (1997). Local likelihood and local partial likelihood in hazard regression. *The Annals of Statistics*, 25(4), 1661 – 1690.

Fan, J., & Li, R. (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American statistical Association*, 96(456), 1348–1360.

Fan, J., & Li, R. (2002). Variable selection for cox’s proportional hazards model and frailty model. *The Annals of Statistics*, 30(1), 74–99.

Faraggi, D., & Simon, R. (1995). A neural network model for survival data. *Statistics in medicine*, 14(1), 73–82.

Faraggi, D., & Simon, R. (1998). Bayesian variable selection method for censored survival data. *Biometrics*, pages 1475–1485.

- Han, K., Wang, Y., Zhang, C., Li, C., & Xu, C. (2018). Autoencoder inspired unsupervised feature selection. In *2018 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pages 2941–2945. IEEE.
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778.
- Ibrahim, J. G., Chen, M.-H., & MacEachern, S. N. (1999). Bayesian variable selection for proportional hazards models. *Canadian Journal of Statistics*, 27(4), 701–717.
- Ishwaran, H., Kogalur, U. B., Blackstone, E. H., & Lauer, M. S. (2008). Random survival forests. *The annals of applied statistics*, 2(3), 841–860.
- Jiménez-Luna, J., Grisoni, F., and Schneider, G. (2020). Drug discovery with explainable artificial intelligence. *Nature Machine Intelligence*, 2(10), 573–584.
- Katzman, J. L., Shaham, U., Cloninger, A., Bates, J., Jiang, T., & Kluger, Y. (2018). Deepsurv: personalized treatment recommender system using a cox proportional hazards deep neural network. *BMC medical research methodology*, 18(1), 1–12.
- Kvamme, H., Borgan, Ø., & Scheel, I. (2019). Time-to-event prediction with neural networks and cox regression. *arXiv preprint arXiv:1907.00825*.
- Lee, C., Zame, W., Yoon, J., & Van Der Schaar, M. (2018). Deephit: A deep learning approach to survival analysis with competing risks. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32.
- Lemhadri, I., Ruan, F., Abraham, L., & Tibshirani, R. (2021). Lassonet: A neural network with feature sparsity. *Journal of Machine Learning Research*, 22(127), 1–29.
- Li, K., Wang, F., Liu, R., Yang, F., & Shang, Z. (2021). Calibrating multi-dimensional complex ode from noisy data via deep neural networks. *arXiv preprint arXiv:2106.03591*.
- Li, K., Wang, F., & Yang, L. (2022). Deep feature screening: Feature selection for ultra high-dimensional data via deep neural networks. *arXiv preprint arXiv:2204.01682*.
- Liu, B., Wei, Y., Zhang, Y., & Yang, Q. (2017). Deep neural networks for high dimension, low sample size data. In *IJCAI*, pages 2287–2293.
- Luck, M., Sylvain, T., Cardinal, H., Lodi, A., & Bengio, Y. (2017). Deep learning for patient-specific kidney graft survival analysis. *arXiv preprint arXiv:1705.10245*.
- Manduchi, L., Marcinkevičs, R., Massi, M. C., Weikert, T., Sauter, A., Gotta, V., Müller, T., Vasella, F., Neidert, M. C., Pfister, M., et al. (2021). A deep variational approach to clustering survival data. *arXiv preprint arXiv:2106.05763*.
- Mirzaei, A., Pourahmadi, V., Soltani, M., & Sheikhzadeh, H. (2020). Deep feature selection using a teacher-student network. *Neurocomputing*, 383, 396–408.
- Nagpal, C., Li, X., & Dubrawski, A. (2021a). Deep survival machines: Fully parametric survival regression and representation learning for censored data with competing risks. *IEEE Journal of Biomedical and Health Informatics*, 25(8), 3163–3175.
- Nagpal, C., Yadlowsky, S., Rostamzadeh, N., & Heller, K. (2021b). Deep cox mixtures for survival regression. In *Machine Learning for Healthcare Conference*, pages 674–708. PMLR.
- Ren, K., Qin, J., Zheng, L., Yang, Z., Zhang, W., Qiu, L., & Yu, Y. (2019). Deep recurrent survival analysis. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 4798–4805.
- Rosenwald, A., Wright, G., Chan, W. C., Connors, J. M., Campo, E., Fisher, R. I., Gascoyne, R. D., Muller-Hermelink, H. K., Smeland, E. B., Giltneane, J. M., et al. (2002). The use of molecular profiling to predict survival after chemotherapy for diffuse large-b-cell lymphoma. *New England Journal of Medicine*, 346(25), 1937–1947.
- Saeys, Y., Inza, I., & Larranaga, P. (2007). A review of feature selection techniques in bioinformatics. *bioinformatics*, 23(19), 2507–2517.
- Sha, N., Tadesse, M. G., & Vannucci, M. (2006). Bayesian variable selection for the analysis of microarray data with censored outcomes. *Bioinformatics*, 22(18), 2262–2268.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1), 267–288.
- Tibshirani, R. (1997). The lasso method for variable selection in the cox model. *Statistics in medicine*, 16(4), 385–395.
- Yi, F., Tang, N., & Sun, J. (2020). Simultaneous variable selection and estimation for joint models of longitudinal and failure time data with interval censoring. *Biometrics*.
- Zhang, H. H., & Lu, W. (2007). Adaptive lasso for cox’s proportional hazards model. *Biometrika*, 94(3), 691–703.
- Zhao, H., Wu, Q., Li, G., & Sun, J. (2019). Simultaneous estimation and variable selection for interval-censored data with broken adaptive ridge regression. *Journal of the American Statistical Association*.
- Zhong, Q., Mueller, J. W., & Wang, J.-L. (2021). Deep extended hazard models for survival analysis. *Advances in Neural Information Processing Systems*, 34.
- Zhu, X., Yao, J., Zhu, F., & Huang, J. (2017). Wsisa: Making survival prediction from whole slide histopathological images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7234–7242.

Zou, H. (2006). The adaptive lasso and its oracle properties. *Journal of the American statistical association*, 101(476), 1418–1429.

### **Copyrights**

Copyright for this article is retained by the author(s), with first publication rights granted to the journal.

This is an open-access article distributed under the terms and conditions of the Creative Commons Attribution license (<http://creativecommons.org/licenses/by/4.0/>).