

Genetics and population analysis

Tutorial on survival modeling with applications to omics data

Zhi Zhao ^{1,2,*}, John Zobolas ^{1,2}, Manuela Zucknick ^{1,3,†}, Tero Aittokallio ^{1,2,4,†}

¹Oslo Centre for Biostatistics and Epidemiology (OCBE), Department of Biostatistics, Faculty of Medicine, University of Oslo, Oslo 0372, Norway

²Department of Cancer Genetics, Institute for Cancer Research, Oslo University Hospital, Oslo 0310, Norway

³Oslo Centre for Biostatistics and Epidemiology (OCBE), Research Support Services, Oslo University Hospital, Oslo 0372, Norway

⁴Institute for Molecular Medicine Finland (FIMM), HiLIFE, University of Helsinki, Helsinki FI-00014, Finland

*Corresponding author. Oslo Centre for Biostatistics and Epidemiology (OCBE), Department of Biostatistics, University of Oslo, P.O.Box 1122, Blindern, Oslo 0317, Norway. E-mail: zhi.zhao@medisin.uio.no (Z.Z.)

†These authors equally contributed to the work.

Associate Editor: Jonathan Wren

Abstract

Motivation: Identification of genomic, molecular and clinical markers prognostic of patient survival is important for developing personalized disease prevention, diagnostic and treatment approaches. Modern omics technologies have made it possible to investigate the prognostic impact of markers at multiple molecular levels, including genomics, epigenomics, transcriptomics, proteomics and metabolomics, and how these potential risk factors complement clinical characterization of patient outcomes for survival prognosis. However, the massive sizes of the omics datasets, along with their correlation structures, pose challenges for studying relationships between the molecular information and patients' survival outcomes.

Results: We present a general workflow for survival analysis that is applicable to high-dimensional omics data as inputs when identifying survival-associated features and validating survival models. In particular, we focus on the commonly used Cox-type penalized regressions and hierarchical Bayesian models for feature selection in survival analysis, which are especially useful for high-dimensional data, but the framework is applicable more generally.

Availability and implementation: A step-by-step R tutorial using The Cancer Genome Atlas survival and omics data for the execution and evaluation of survival models has been made available at <https://ocbe-uio.github.io/survomics>.

1 Introduction

Personalized medicine improves patient diagnosis and treatment by making use of patient-specific genomic and molecular markers that are indicative of disease development. Time to an event of interest (e.g. time to death or disease progression) is a widely used end point and patient outcome for many diseases, and therefore it has become popular to identify genomic, molecular and clinical markers for survival or progression prediction of patients suffering from complex diseases such as cancer. In this tutorial, we will only consider so-called right-censored survival data, where a patient has been followed for a certain time period and the event of interest is either observed in this time period or might occur at a later (as yet unobserved) time point. Right-censored survival data include both a time and the status of each patient at that time as joint outcomes, where time is a continuous variable and status is a binary variable indicating whether the event of interest has been observed up to the given time or not; in the latter case, we refer to this observation as censored at the observed time. See [Table 1](#) for an example illustration, and more details of the survival data in Section 2. When using the status label as an outcome in an ordinary logistic regression,

the regression coefficients will become increasingly uncertain and less reliable with increasing follow-up time ([Green and Symons 1983](#)). When using the observed time (or its transformation, e.g. logarithm of time) as an outcome in an ordinary linear regression, the presence of censored observations (i.e. patients still alive by the end of follow-up period) causes considerable difficulties for assessing the accuracy of predictions ([Henderson 1995](#)). For example, [Table 1](#) shows an example where the proportion of patients surviving past 10 years is $1/5 = 20\%$ based on the observed data, but the (unobserved) factual proportion surviving past 10 years is $3/5 = 60\%$.

The availability of multiple types of genomic and molecular data poses great opportunities but also further challenges for building effective statistical models to identify biomarkers that are prognostic of patient survival. For example, omics profiles, such as those from mRNA expression, DNA copy number, single-point and other genetic mutations, may be available from the same patient, and these high-dimensional data come with intra- and inter-dataset correlations, heterogeneous measurement scales, missing values, technical variability, and other background noise ([Hasin *et al.* 2017](#)). [Rahnenführer *et al.* \(2023\)](#) provided a general guideline for

Table 1. An example of right-censored time-to-event (survival) outcomes, illustrated in an example where the outcome of interest is the time from disease diagnosis to death.^a

Patient ID	Observed time (years)	Status at observed time	Factual time to death, possibly unobserved (years)
1001	11	Censored	20
1002	4	Dead	4
1003	5	Censored	12
1004	9	Dead	9
1005	1	Censored	11

^a If a patient was alive at the last observed time point, the measurement is censored at that time point.

high-dimensional omics and electronic health records data analysis, and discussed statistical challenges and opportunities for survival modeling. [Bøvelstad *et al.* \(2007\)](#) and [Bøvelstad *et al.* \(2009\)](#) compared various machine learning methods and showed improved survival prediction performance by coefficient shrinkage methods that combine several data sources, in their case clinical and gene expression data. Another recent study by [Bommert *et al.* \(2022\)](#) performed a benchmark of filter methods for feature selection in high-dimensional gene expression data for survival prediction, and recommended using the simple variance filter before fitting a ℓ_2 -regularized Cox regression model ([Simon *et al.* 2011](#)) for accurate survival prediction and stable feature selection. [Vinga \(2021\)](#) reviewed structured penalized regressions for analyzing high-dimensional omics data for survival prediction or evaluation. Multiple studies have shown that it is possible to further improve the prediction accuracy and feature selection by considering more complex structures, such as biological pathways, or by identifying significant features among functional relationships between the omics features ([Chekouo *et al.* 2015](#), [Kundu *et al.* 2018](#), [Wang and Liu 2020](#), [Madjar *et al.* 2021](#)).

In this tutorial, we describe a general workflow for survival analysis with omics data, as well as review the commonly used statistical methods for feature selection and survival prediction; importantly, we provide a step-by-step R tutorial using publicly available omics data from The Cancer Genome Atlas (TCGA) project (<http://cancergenome.nih.gov>). In this example dataset, the overall survival time, demographic and gene expression data from primary invasive breast cancer patients in TCGA (TCGA-BRCA) were retrieved from the Genomic Data Commons Data Portal data release 32.0–36.0. Compared to the previous reviews and benchmarks of survival models in bioinformatic applications, this tutorial provides a complete workflow ranging from data preparation to final calibrated models, with a particular focus on building survival models using high-dimensional omics data, and as such covers both the commonly used penalized regressions and Bayesian models for survival analysis with high-dimensional and generally noisy datasets. For this tutorial, we assume that readers have the knowledge of basic statistical methods. Terms beyond basic statistics are explained in corresponding text.

2 Data categories

We use omics data and overall survival outcomes from cancer patients as an example in this tutorial, but the methods are

applicable also to other diseases with similar data types in personalized medicine applications. The ultimate goal of personalized medicine is to identify patient-specific risk factors to guide disease prevention, diagnostic and treatment strategies. The identification of potential risk factors for cancer patients often considers clinical, demographic, genomic and molecular information, and their associations with the patients' time-to-event data (i.e. survival).

2.1 Time-to-event data

Time-to-event or survival data contain the event of interest (e.g. death is the event assumed in this section), together with the time from the beginning of a study period either to the occurrence of the event, end of the study, or patient loss to follow-up (i.e. right censoring; discussions of right, left and interval censoring can be found in [Leung *et al.* 1997](#)). [Figure 1c](#) shows the survival or right-censored times since cancer diagnosis of patients. The exact survival time of a patient may be not observed due to censoring. Therefore, a patient has two outcome indices: censoring indicator δ (also called status) and observed time $\tilde{T} = \min\{T, C\}$, where T is the exact survival time and C is the censoring time. Indicator δ can also be denoted as $\delta = \mathbb{1}\{T \leq C\}$, where $\mathbb{1}\{\cdot\}$ is an indicator function. To characterize the survival time of a patient, we can use survival function $S(t) = \mathbb{P}\{T > t\}$, which gives the probability of the patient's survival beyond time t . Another useful quantity is the hazard function

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

which is the instantaneous probability of the patient's death at time t conditional on the patient having survived up to that time point.

2.2 Clinical and demographic data

There are multiple sources of patient-level information that can be explored to identify risk factors for cancer patients. One can start with routinely collected and commonly used patient data, such as clinical and demographic variables. For example, an older male patient with a low body mass index (BMI) has a relatively high risk of gastric cancer ([Nam *et al.* 2022](#)). [Table 2](#) illustrates selected clinical and demographic variables often available for cancer patients. Clinical and demographic variables are considered important sources of information for predicting survival and are often used to build reference models for omics-based prognostic models ([Herrmann *et al.* 2021](#)).

2.3 Omics data

Thanks to the rapid development of modern molecular biotechnology, large amounts of human genomic and molecular data have become available from many patient profiling projects. These projects often collect multiple levels of molecular information such as genomics data for DNA variation, transcriptomics data for mRNA expression, proteomics data for protein abundance and metabolomics data for metabolite processes, as illustrated in [Fig. 1a](#). Among the multiple omics levels, metabolomics is the closest to observable phenotypes, such as tumor growth and proliferation ([Cairns *et al.* 2011](#)). To deeply understand the molecular biology of tumor development, multiple levels of omics data may deliver novel insights into the circuits of molecular interactions that

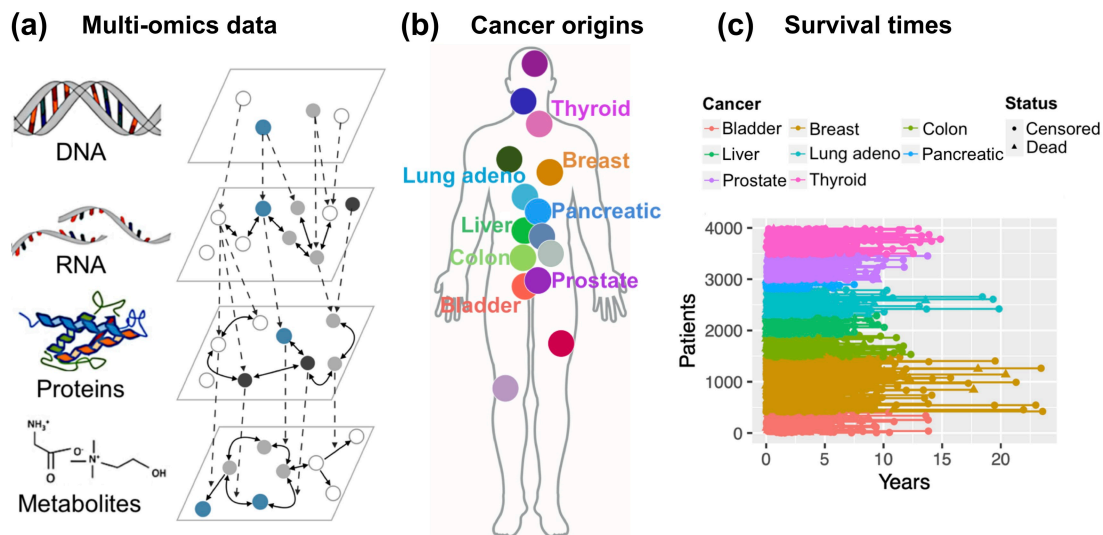


Figure 1. Pan-cancer survival data and omics signatures. (a) Multiple omics layers [modified from Haukaas *et al.* (2017) and Jendoubi (2021)]. The network illustrates intra- and cross-layered biological features or molecules (e.g. DNA methylation, mRNA, proteins, metabolites). (b) Different origins of tumors represented by the TCGA pan-cancer patient cohorts. (c) Overall survival times of cancer patients from TCGA. Novel methods are needed to model the high-dimensional multi-omics data and leverage information from heterogeneous cancer cohorts for improved survival prognosis and biomarker identification.

Table 2. Examples of clinical and demographic variables.

Variable	Data type
Sex	Binary
Body mass index (BMI, kg/m ²)	Continuous
Ethnicity	Nominal
Age at first diagnosis in years	Integer
Pathological stage	Ordinal
Therapy type (e.g. chemo-, hormone, immuno-therapy)	Nominal

underlie the disease initiation and progression (Tarazona *et al.* 2021).

DNA-level omics data often include single nucleotide polymorphisms (SNP), DNA methylation, and somatic copy number variation, see illustration of these data types in Table 3. Each SNP feature can be coded as $\{0, 1, 2\}$, according to the number of minor alleles at the given locus in each individual, i.e. AA, AG, GC, or TT. DNA methylation reveals methyl groups added to the DNA molecule, which can be quantified as a β -value $\beta = \frac{M}{M+U+a} \in [0, 1]$, where M and U are the fluorescence intensities of the methylated and unmethylated DNA at a CpG locus, and the offset a is often set to 100 to stabilize the β -value when M and U are small. Somatic copy number variation measures the number of repeated sections of the tumor genome. Table 3 shows the assumed distributions for the downstream statistical modeling. For example, DNA methylation β -value is a proportion, often either close to 1 or 0, which can be characterized by a mixture of two beta distributions.

RNA-level omics data usually include messenger RNA (mRNA) expression and microRNA (miRNA) expression. Traditionally, DNA microarrays were used to measure the expression levels of DNA sequences called probes, which acted as a proxy for the amount of reads representing a genomic feature of interest. The reads of microarray expression level can be characterized by a negative binomial distribution directly, or a Gaussian distribution after \log_2 transformation (see Table 3). Nowadays, RNA sequencing (RNA-seq) has replaced DNA microarrays, since it allows for the sequencing

of the whole transcriptome, while DNA microarrays only profile predefined transcripts or gene probes. miRNAs are small noncoding regulatory RNAs that play an important role in regulating gene expression and are highly evolutionary conserved in mammals (Bartel 2018).

Protein-level omics data usually originates from mass spectrometry (MS)-based proteomics profiling of a particular cell, tissue, organ, which can detect and measure abundance levels of entire or phosphorylated proteins. In contrast to global proteomics with MS, TCGA consortium has produced more targeted proteomics profiles, involving a set of 180–250 protein features using reverse-phase protein arrays (RPPA) (Akbani *et al.* 2014). In contrast, The Clinical Proteomic Tumor Analysis Consortium (CPTAC) profiling has produced >10 000 protein features using MS technology (Edwards *et al.* 2015). Protein expression data can often be considered approximately Gaussian distributed after \log_2 transformation, depending on the data-generating platform. In global proteomics, there are often sample-specific missing values due to detection limits of protein quantification.

Metabolite-level omics data has similar statistical properties to proteomics data, and metabolomics profiling is usually done with MS-based technologies, enabling the detection, and quantification of many thousands of metabolic features simultaneously. Nuclear magnetic resonance (NMR) spectroscopy is the other main analytical technology to profile metabolic processes. An important limitation of NMR spectroscopy is its relatively low sensitivity, which may lead to relatively few detected metabolites. Metabolic concentrations often follow logarithmic Gaussian distribution. Similar to large-scale proteomics, missing values due to detection limits should be treated differently than missing values due to measurement artefacts, which are more frequent in metabolite-level omics data (Sun and Xia 2023).

Single-cell sequencing is becoming increasingly more prevalent in many profiling studies. The modern single-cell omics technologies can produce multiple levels of omics data derived from the same samples, such as transcriptomics and chromatin accessibility. The single-cell data types are similar

Table 3. Illustration of omics data and distribution assumptions for data generated by commonly used high-throughput technologies.

Patient ID	SNP	Methylation (β -value)	Copy number variation (number of copies)	Gene expression (reads)	miRNA (reads)	Protein (intensity/concentration)	Metabolite (intensity/concentration)
1001	1	0.2	0	5	5	0.07	0.07
1002	0	0.11	5	2	2	0.1	0.1
1003	1	0.95	10	0	0	9.6	9.6
1004	2	0.5	4	30	30	2.8	2.8
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
Technology	SNP array	Infinium BeadChip	SNP array (Gistic2)	RNA-Seq	RNA-Seq	Mass spectrometry	Mass spectrometry
Distribution assumptions	Ordinal	Beta mixture	Negative binomial log ₂ ratio: Gaussian	Negative binomial log ₂ scale: Gaussian	Negative binomial log ₂ scale: Gaussian	log-Gaussian log ₂ scale: Gaussian	log-Gaussian log ₂ scale: Gaussian

to the illustration in Table 3, but each measurement originates from the level of individual cells. This article focuses mainly on survival analysis with single bulk omics data.

2.4 Missing data

Missing values are often observed in many types of high-dimensional omics data due to various experimental reasons (Aittokallio 2009, Kong et al. 2022). For example, mRNA transcriptomics data from microarrays have 1%–10% missing values affecting up to 95% of genes due to corruption of image, scratches on the slides, poor hybridization, inadequate resolution, fabrication errors (de Brevern et al. 2004, Tuikkala et al. 2005); MS-based metabolomics data have 10%–20% missing values affecting up to 80% of variables due to lack of peak identification by chromatogram, limitation of computational detection, measurement error, and deconvolution errors (Hrydziuszko and Viant 2012, Sun and Xia 2023). The aforementioned technical reasons can lead to missing data that is either missing at random (MAR) or missing not at random (MNAR). When dealing with MNAR data, traditional imputation methods like multiple imputation may introduce bias. It is thus recommended to remove omics features which have large proportion (e.g. 50%) missingness over patients, and then apply imputation methods (e.g. k -nearest neighbor imputation) for the rest of the features with missing values before doing any statistical analysis or modeling. Alternatively, imputation-free methods (e.g. mixture models) that can deal with missing values can be applied directly (Taylor et al. 2022). Single-cell RNA-seq (scRNA-seq) data has a vast number of zeros, so-called gene dropout events, leading highly scarce data matrices. Jiang et al. (2022) discussed the sources of biological and nonbiological zeros in scRNA-seq data and the existing approaches for handling them.

3 Survival analysis with low-dimensional input data

Let us assume we have data $\mathcal{D} = \{(\tilde{T}_i, \delta_i, \mathbf{X}_i) : i = 1, \dots, n\}$ for n patients, where \tilde{T}_i is the observed survival time, δ_i the censoring indicator and \mathbf{X}_i contains p covariates including clinical, demographic and omics features. To estimate a survival function $S(t)$ given the data \mathcal{D} , one needs to keep track both of the number of patients at risk and those who left the study at any time point (here we only consider the case of right censoring and assume no delayed entry). At time t there are $Y(t) = \sum_{i=1}^n Y_i(t)$ patients at risk, where $Y_i(t) =$

$\mathbb{1}\{\tilde{T}_i \geq t\}$ is the indicator that patient i is at risk. The non-parametric Kaplan–Meier (KM) estimator (Kaplan and Meier 1958) uses the multiplication rule for conditional probabilities to obtain an estimation of the survival function $\hat{S}(t) = \prod_{T_k \leq t}^K \left(1 - \frac{d_k}{Y(T_k)}\right)$, where all events happen (e.g. patients

die) at K distinct times, $T_1 < T_2 < \dots < T_K$ ($K \leq n$) and there are $d_k \geq 1$ events happened at the time T_k . If no two events happen at the same time point, $d_k = 1$ and $k = 1, \dots, n$. The KM estimator gives an estimate of the marginal survival function, i.e. when you disregard the information from the covariates.

Figure 2a shows the KM curve for TCGA-BRCA primary breast cancer patients data. Some basic statistics can be revealed from the survival curve. For example, the estimated median survival time, i.e. the time when the survival probability is 50%, of all the breast cancer patients is 10.8 years (dashed line in Fig. 2a), and 1-, 5-, and 10-year survival probabilities are 0.988, 0.853, and 0.658, respectively. A log-rank test (Peto and Peto 1972) can be used to test whether two groups of patients [e.g. with treatment (pharmaceutical/radiation therapy) or nontreatment] have the same (null hypothesis) or different survival functions (alternative hypothesis), and provide a corresponding P -value (see Fig. 2b). The log-rank test can also be used to compare the survival probabilities of any subgroups of patients based on other categorical variables.

In the case where multiple clinical, demographic or omics features are available, one can explore each variable's association with survival outcomes separately. For a categorical variable, KM curves and log-rank tests can be used to investigate whether there is a difference between multiple survival curves categorized by the variable. For a continuous variable X , the semi-parametric Cox proportional hazards model (Cox model, Cox 1972) is often used:

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

where $h_0(t)$ is the baseline hazard function and is left unspecified. As the name ‘proportional hazards’ implies, the Cox-model estimated hazard functions of two individuals (with different values of the covariate X) are indeed proportional, because $h_0(t)$ does not depend on X and is thus assumed to be the same for all individuals. The functional form (1) describes the log-linear relationship between variable X and the hazard $h(t|X)$ at any given timepoint t . It may be difficult

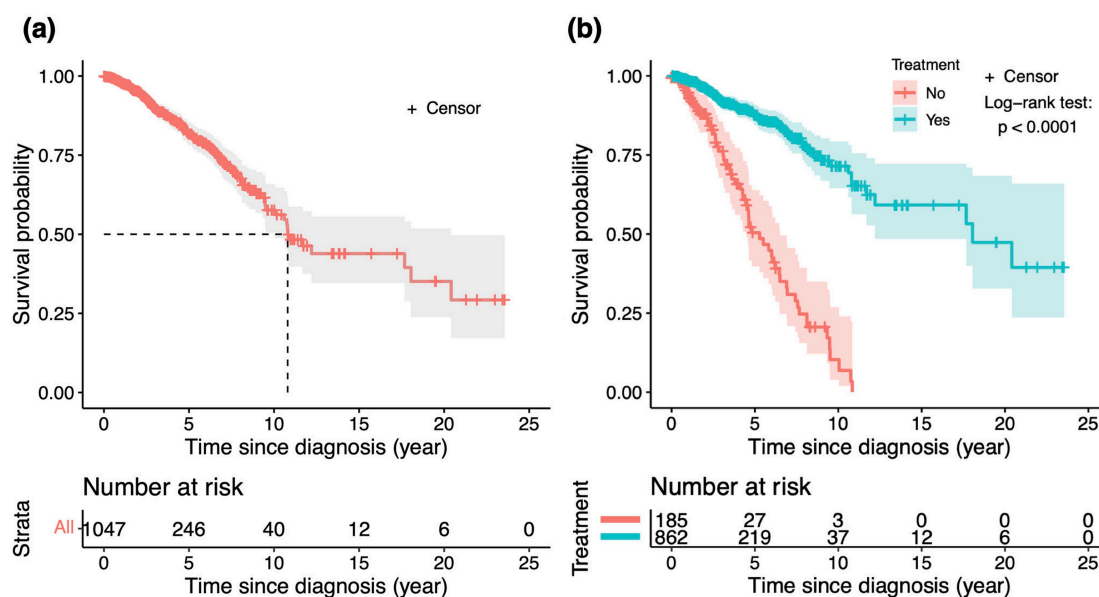


Figure 2. Kaplan–Meier curves of TCGA-BRCA data. (a) KM curve of all the TCGA-BRCA patients’ survival data. (b) KM curves of the TCGA-BRCA patients’ survival data grouped by treatment (i.e. pharmaceutical/radiation therapy) or nontreatment. The log-rank test is used to compare the two survival distributions corresponding to the two groups of patients.

to satisfy the log-linear relationship based on the original scale of some omics features, e.g. gene expression data from a DNA microarray study; in those cases, the use of \log_2 -transformation of the data (Table 3) can be helpful. The Cox model can also be used to investigate the risk of a categorical variable, provided the above assumptions are satisfied, which provides an effect estimate (hazard ratio, HR), in addition to the associated P -value.

It is often of interest to gain insights into the multiple factors and their cooperation for the survival outcomes. Multivariable statistical analysis plays an important role in such multi-modal survival modeling. The univariate Cox model (1) can be straightforwardly generalized by including multiple clinical, demographic or omics variables of interest, as long as the total number of covariates is much smaller than the number of samples; this often requires the use of variable or feature selection methods. Heinze *et al.* (2018) provides pragmatic recommendations for practitioners of multivariable analysis with variable selection methods in low-dimensional modeling problems.

4 Survival analysis with high-dimensional input data

Since some omics data contain high-number of variables (e.g. RNA-seq with ca. 60 000 transcriptomic features and DNA methylation with ca. 450 000 features), there is a need to reduce the computational and modeling burden in multivariable analyses. One heuristics approach to pre-select a subset of features is to include only omics features at a specific statistical significance level when fitting a univariate Cox model (1); the pre-specified significance level would usually be higher than the commonly used 0.05 threshold, e.g. 0.1 or 0.2, to avoid losing important features. However, this univariate approach focuses on features that are independently associated with the outcome and might miss variables that are important in combination with other features (Okser *et al.* 2013).

Another simple approach is to pre-select omics features by variance, since larger variability across patients usually implies higher biological information, or at least predictive signal. One can for example pre-select omics features explaining 50% of the total cumulative variance (Zhao and Zucknick 2020). Such unsupervised feature pre-selection is a recommended method to reduce dimensionality when dealing with hundreds or thousands of omics features, with the aim to improve the stability of final feature selection. For example, Bommert *et al.* (2022) showed that the simple variance filter was the best method among all considered filter methods in terms of the predictive accuracy, run time and the feature selection stability in their benchmark study. Zhao *et al.* (2024) showed that the stability of final feature subset critically depends on the pre-selected feature set when using a standard Bayesian stochastic search variable selection method (see Section 4.3), and their proposed method that used known biological relationships between omics features lead to a more stable feature selection and slightly improved outcome prediction.

When drawing conclusions on survival differences solely from the univariate Cox model (1), it is important to adjust the P -values of risk features for multiple comparisons to control false positives globally, e.g. by controlling the family-wise error rate (FWER) or false discovery rate (FDR). The problem of multiple testing is beyond the scope of this tutorial, but the interested readers are referred to a recent review article (Korthauer *et al.* 2019). We note also that univariate analysis does not consider any relationships between multiple omics features, e.g. potential confounders for survival outcomes (Clark *et al.* 2003). To avoid making seriously misleading conclusions in such cases, it is necessary to perform multivariable survival analysis (Bradburn *et al.* 2003).

Omics data can have hundreds of thousands of variables measured at various molecular levels, which greatly challenges the classical multivariable regression models for time-to-event endpoints, since the number of variables is often much larger than the number of patients (i.e. $p \gg n$). To

proceed with survival analysis, one option is to reduce the dimensionality of the omics features via unsupervised learning, and then investigate the association of the learned low-dimensional variables with survival outcomes (see next section). An alternative approach is to directly use supervised learning methods, such as penalized regressions or sparse Bayesian models (Table 4), which enable the modeling of high-dimensional omics features, and the selection of key important features associated with the survival outcomes.

4.1 Unsupervised learning

Unsupervised learning methods aim to identify hidden patterns or data groupings, and are for example useful when a phenotype (e.g. a cancer type) is to be divided into several subtypes (e.g. to explain heterogeneity among patients). For example, breast cancer has been traditionally categorized into five conceptual molecular classes, originally using pairwise average-linkage cluster analysis of DNA microarray data, to better understand tumor biology and guide clinical decision making (Perou *et al.* 2000). Unsupervised methods learn underlying patterns from unlabeled data by transforming high-dimensional omics features into a lower dimensional space. Principal component analysis (PCA) is a classical multivariate technique that represents high-dimensional features in a low-dimensional space by building orthogonal (uncorrelated) linear combinations of the features that best capture the variance in the high-dimensional data.

Different from the distance-based PCA with linear transformation, nonlinear techniques have recently emerged, such as *t*-stochastic neighbor embedding (*t*-SNE). *t*-SNE uses pairwise similarities of individuals based on Kullback–Leibler divergence to project similarities into a lower dimensional space, ensuring that individuals with similar omics features are close in the generated embedding (van der Maaten and Hinton 2008). The focus of *t*-SNE is on preserving local distances between neighboring data points. Another nonlinear alternative to PCA is UMAP (Uniform Manifold Approximation and Projection), which is a general dimension

reduction method built upon Riemannian geometry (McInnes *et al.* 2018). Compared to other nonlinear methods of dimension reduction such as *t*-SNE, UMAP can sometimes provide better visualization quality in a shorter amount of time, while preserving the global structure of the omics data better.

Unsupervised methods only make use of the input data matrix \mathbf{X} and are agnostic to the survival information. To find out whether a given omics profile (i.e. the full set of omics features, not individual features) is associated with survival outcomes, a straightforward approach is to use a few representative components from PCA, *t*-SNE, or UMAP as covariates in a multivariable Cox model. An alternative is to use semi-supervised methods (Bair and Tibshirani 2004) that combine the clustering procedure and survival modeling together. However, such principal component regression and semi-supervised methods lose interpretability of the individual omics features, since each component is a linear or nonlinear combination of all omics features.

4.2 Supervised learning via penalized regressions

For the purpose of personalized cancer medicine, one is typically interested in identifying risk factor combinations from clinical and omics features. These factors can be targeted (directly or indirectly) via therapeutic strategies or used for diagnostics. Therefore, the objective is to identify a parsimonious set of features linked to survival outcomes by utilizing the wealth of information present in, e.g. the vast amount of available omics data. Penalized Lasso Cox regression (Tibshirani 1997) can be used to select a few relevant omics features by estimating their coefficients as nonzero (the non-relevant features' coefficients are shrunk to zero) via maximizing the penalized partial log-likelihood function of the regression coefficients with ℓ_1 -norm penalty

$$2/n \cdot \ell(\boldsymbol{\beta}|\mathcal{D}) - \lambda \|\boldsymbol{\beta}\|_1 \quad (2)$$

Here, $2/n$ is a scaling factor for convenience, $\mathcal{D} = \{(\bar{T}_i, \delta_i, \mathbf{X}_i) : i = 1, \dots, n\}$, \mathbf{X}_i includes p (omics) features of

Table 4. Cox-type supervised learning methods.

Method	Feature selection via	Grouping effects	Uncertainty quantification	Comment
Penalized regressions:	Penalty			
Lasso Cox (Tibshirani 1997)	ℓ_1 -norm	✗	✗	
Adaptive Lasso Cox (Zhang and Lu 2007)	Weighted ℓ_1 -norm	✗	✗	Less false positives than Lasso
Elastic Net Cox (Simon <i>et al.</i> 2011)	ℓ_1/ℓ_2 -norm	✓	✗	
Group-Lasso Cox (Kim <i>et al.</i> 2012)	ℓ_2 -norm	✓	✗	Independent groups of features selected
Sparse Group-Lasso Cox (Simon <i>et al.</i> 2013)	ℓ_1/ℓ_2 -norm	✓	✗	
SCAD Cox (Fan and Li 2002)	Quadratic spline and symmetric penalty	✓	✗	Selection of relatively large effects
SIS Cox (Fan <i>et al.</i> 2010)	Top-ranked variables and any penalty	✓	✗	Suited to ultra-high dimensions
Bayesian models:	Shrinkage prior			
Lee <i>et al.</i> (2011)	Lasso (Laplace) prior	✗	✓	Selection of posterior mean with a cutoff
Lee <i>et al.</i> (2015)	Elastic Net, group/fused Lasso priors	✓	✓	Selection of posterior mean with a cutoff
Konrath <i>et al.</i> (2013)	Spike-and-slab prior	✗	✓	
Madjar <i>et al.</i> (2021)	Spike-and-slab and MRF priors	✓	✓	
Mu <i>et al.</i> (2021)	Horseshoe prior	✗	✓	Selection of posterior mean with a cutoff

the i th patient, $\lambda \geq 0$ is a tuning parameter to control the overall penalty strength of the coefficients, $\|\boldsymbol{\beta}\|_1 = \sum_{j=1}^p |\beta_j|$, and the partial log-likelihood is $\ell(\boldsymbol{\beta}|\mathcal{D}) = \log \prod_{i=1}^n \left\{ \frac{\exp(\mathbf{X}_i \boldsymbol{\beta})}{\sum_{l \in \mathcal{R}_k} \exp(\mathbf{X}_l \boldsymbol{\beta})} \right\}^{\delta_i}$, where $\mathcal{R}_k = \{l : Y_l(T_k) = 1\}$ is the risk set at time T_k . The Elastic Net Cox model (Simon *et al.* 2011) considers both the Lasso feature selection and the grouping effect of correlated omics features in ridge regression via a combination of the ℓ_1 - and ℓ_2 -norm (i.e. ℓ_1/ℓ_2 -norm) penalty $\lambda \{ \alpha \|\boldsymbol{\beta}\|_1 + \frac{1}{2}(1-\alpha) \|\boldsymbol{\beta}\|_2^2 \}$ (where $\|\boldsymbol{\beta}\|_2^2 = \sum_{j=1}^p |\beta_j|^2$), which can usually improve the prediction performance over the Lasso Cox model. Figure 3 shows an example of Elastic Net Cox model feature selection from gene expression features associated with breast cancer patients' survival. Note that often we wish to include a small set of well-established clinical risk factors in a survival model. Since they are established as important covariates, they can be included in the Lasso or Elastic Net Cox model as mandatory covariates without penalization. Then the penalized partial log-likelihood function becomes

$$\frac{2}{n} \log \prod_{i=1}^n \left\{ \frac{\exp(\mathbf{X}_{0i} \boldsymbol{\beta}_0 + \mathbf{X}_i \boldsymbol{\beta})}{\sum_{l \in \mathcal{R}_k} \exp(\mathbf{X}_{0l} \boldsymbol{\beta}_0 + \mathbf{X}_l \boldsymbol{\beta})} \right\}^{\delta_i} - \text{pen}(\boldsymbol{\beta}), \quad (3)$$

where $\boldsymbol{\beta}_0$ are coefficients corresponding to the i th individual's mandatory covariates \mathbf{X}_{0i} , $\text{pen}(\boldsymbol{\beta})$ is a ℓ_1 - or ℓ_1/ℓ_2 -norm penalty for feature selection of omics features \mathbf{X}_i . De Bin *et al.* (2014) investigated more strategies to combine a low-dimensional set of well-established clinical factors and high-dimensional omics features into a global prediction model.

There are many alternative penalties that will achieve feature selection which have been applied to Cox proportional hazards regression for survival outcomes, such as the Adaptive Lasso Cox model that incorporates different penalties for different coefficients to retain important variables (Zhang and Lu 2007), Group-Lasso that performs group selection on (predefined) groups of variables (Kim *et al.* 2012),




Figure 3. Coefficient trace plot of an Elastic Net Cox model for overall survival prognosis of breast cancer patients from TCGA based on transcriptomic data and mandatory demographic variables age and ethnicity. The y -axis shows the magnitude of each feature's coefficient given the strength of penalization displayed on the x -axis (from left to right the penalization decreases). The vertical line indicates the optimal $\lambda = 0.032$ (maximizes the partial likelihood via cross-validation) and its corresponding selected feature names are listed on the right side. Note that the demographic variables age and ethnicity were not penalized, so that their coefficient paths did not start from zero in the figure.

Sparse Group-Lasso that introduces sparse effects both on a group and within group level (Simon *et al.* 2013), the smoothly clipped absolute deviation (SCAD) Cox model that overcomes substantially biased estimates for large coefficients in ultra-sparse models (Fan and Li 2002), sure independence screening (SIS) procedure in combination with Cox model that speeds-up the feature selection dramatically and can also improve the accuracy of estimation when dimensionality becomes ultra-high, i.e. $\log(p) = \mathcal{O}(n^\zeta)$ for some $\zeta > 0$ (Fan *et al.* 2010). However, all these penalized Cox models do not directly provide uncertainty of feature selection or survival prediction. One empirical way for uncertainty quantification is through additional resampling-based methods, see Section 5 for more details.

4.3 Supervised learning via Bayesian priors

Bayesian inference is an appealing approach for survival analysis due to its ability to provide straightforward uncertainty quantification (e.g. credible intervals) of parameters and survival probabilities. For instance, Lee *et al.* (2011) proposed a Bayesian version of the Lasso Cox (Bayesian Lasso Cox) model that provides credible intervals of coefficients fairly straightforward (see Fig. 4), but it is not easy to derive confidence intervals of coefficients in a Lasso-type model.

The fundamental theorem of Bayesian methods is Bayes' rule. Let $\boldsymbol{\beta}$ be the parameters of interest, e.g. gene effects, external to the data \mathcal{D} . To estimate the parameters $\boldsymbol{\beta}$ given the data information, one can use Bayes' rule to obtain the conditional (posterior) distribution of $\boldsymbol{\beta}$:

$$f(\boldsymbol{\beta}|\mathcal{D}) = \frac{f(\mathcal{D}|\boldsymbol{\beta})f(\boldsymbol{\beta})}{f(\mathcal{D})} \propto f(\mathcal{D}|\boldsymbol{\beta})f(\boldsymbol{\beta}),$$

where $f(\mathcal{D})$ is a normalization constant that can be neglected in inference since the data are already observed, $f(\mathcal{D}|\boldsymbol{\beta})$ is the likelihood of the data viewed as a function of the parameters of a statistical model and $f(\boldsymbol{\beta})$ is the prior distribution of $\boldsymbol{\beta}$. The prior distribution can be chosen either based on historical data from past similar studies or from popular (non)informative priors (Ibrahim *et al.* 2001). The estimation of $\boldsymbol{\beta}$ is to maximize the posterior $f(\boldsymbol{\beta}|\mathcal{D})$ or $\log f(\boldsymbol{\beta}|\mathcal{D})$, which is equivalent to maximize the sum of the log-likelihood and the log prior, i.e. $\log f(\boldsymbol{\beta}|\mathcal{D}) = \log f(\mathcal{D}|\boldsymbol{\beta}) + \log f(\boldsymbol{\beta})$, which takes into account both the observed data information and the prior information in an optimal way.

The Bayesian version of the Lasso Cox model can have a log posterior similar to the frequentist penalized partial log-likelihood function (2), if we assign independent double exponential (also known as Laplace, Fig. 5a) prior, $f(\boldsymbol{\beta}) = \prod_{j=1}^p \frac{\lambda}{2} \exp\{-\lambda|\beta_j|\}$ with a scale parameter $\lambda > 0$, for all the coefficients, i.e.

$$\log f(\boldsymbol{\beta}|\mathcal{D}, \lambda) = \ell^*(\boldsymbol{\beta}|\mathcal{D}) - \lambda \|\boldsymbol{\beta}\|_1 + C, \quad (4)$$

where $\ell^*(\boldsymbol{\beta}|\mathcal{D})$ is the full log-likelihood function i.e. $\log(\int_0^t h_0(s) ds) + \ell(\boldsymbol{\beta}|\mathcal{D})$, C is a normalization constant independent of $\boldsymbol{\beta}$ and the ℓ_1 -norm penalty tends to choose only a few nonzero coefficients. Markov chain Monte Carlo (MCMC) sampling can be performed for posterior inference of $\boldsymbol{\beta}$. Note that instead of using the partial log-likelihood $\ell(\boldsymbol{\beta}|\mathcal{D})$ in (4), a full log-likelihood function is used, which includes the baseline hazard function. This can be achieved, e.g. by assigning a stochastic process prior, e.g. a gamma process, for the cumulative baseline hazard function. More




Figure 4. Point estimation and uncertainty quantification of regression coefficients using Bayesian Lasso Cox model with Laplace prior for overall survival prognosis of breast cancer patients from TCGA based on transcriptomic data and mandatory demographic variables age and ethnicity. Solid dots indicate the posterior mean over 20000 MCMC iterations excluding burn-in period, and horizontal lines show the corresponding 95% credible intervals.

details about the prior setup and inference can be found in [Lee *et al.* \(2011\)](#). [Lee *et al.* \(2015\)](#) extended the Laplace prior to Elastic Net prior, fused Lasso prior and group Lasso prior, which are often more suitable for correlated omics features in survival analysis. But [Lee *et al.* \(2011, 2015\)](#) assign the same shrinkage priors to all covariates indiscriminately, see [Zucknick *et al.* \(2015\)](#) for an extended Bayesian Lasso Cox model which permits the use of mandatory covariates.

Note that the Bayesian version of the Lasso with Laplace-type priors in practice do not result in automatic feature selection, because only the posterior modes of the coefficients are equivalent to the frequentist Lasso solution, while in Bayesian inference one usually focuses on the posterior means as point estimates. As an alternative, a particular omics feature can be excluded if the estimated credible interval of the corresponding coefficient covers zero.

Stochastic search variable selection (SSVS) is an alternative approach to identify important covariates ([George and McCulloch 1993](#), [Konrath *et al.* 2013](#)). SSVS uses independent spike-and-slab priors for regression coefficients, e.g.

$$\beta_j | \gamma_j, \tau_j^2 \sim \gamma_j \mathcal{N}(0, \tau_j^2) + (1 - \gamma_j) \delta_0(\beta_j), \quad (5)$$

where γ_j ($j = 1, \dots, p$) is a latent variable (which can have a

Bernoulli prior with a fixed probability π) for feature selection indicating $\beta_j \neq 0$ if $\gamma_j = 1$ and $\beta_j = 0$ if $\gamma_j = 0$, τ_j^2 is an additional shrinkage parameter which can be assigned with an additional prior (e.g. exponential or inverse gamma prior), and $\delta_0(\cdot)$ is the Dirac delta function. [Figure 5b](#) shows the two components of the spike-and-slab mixture distribution. [Madjar *et al.* \(2021\)](#) proposed graph-structured feature selection priors for Cox model by assigning a Markov random field prior on the latent variables in the spike-and-slab priors (5), in which the graph helps to identify pathways of functionally related genes or proteins that are simultaneously prognostic in different patient cohorts. Formulation (5) implies independence between the priors of the individual β_j in the slab component. In contrast, a variant of the spike-and-slab prior has a g-prior slab ([Zellner 1986](#), [Held *et al.* 2016](#))

$$\beta_\gamma | \gamma \sim \mathcal{N}(0, g \mathcal{I}_{\beta_\gamma, \beta_\gamma}^{-1}),$$

where $\beta_\gamma = \{\beta_j : \gamma_j = 1, j = 1, \dots, p\}$, $\gamma = \{\gamma_j : j = 1, \dots, p\}$, g is either a scalar estimated by Empirical Bayes or assigned with additional prior, and $\mathcal{I}_{\beta_\gamma, \beta_\gamma}^{-1}$ is the expected Fisher information for β_γ .

Another popular shrinkage prior is the horseshoe prior, a continuous and global-local shrinkage prior, in which the global parameters allow for sharing information across omics features and the local parameters allow for adjustments at the individual omics feature level ([Carvalho *et al.* 2009](#), [Mu *et al.* 2021](#)). In a similar setup to the Cox model with spike-and-slab priors in (5), a horseshoe prior for the regression coefficient is

$$\beta_j | \lambda_j^2, \tau^2 \sim \mathcal{N}(0, \lambda_j^2 \tau^2), \quad \lambda_j \sim C^+(0, 1), \quad \tau \sim C^+(0, 1),$$

where the local parameter λ_j and global parameter τ are both half-Cauchy distributed $C^+(\cdot, \cdot)$. With the horseshoe prior, the posterior mean of β_j will be shrunk by a weight $\kappa_j = \frac{1}{1 + \lambda_j^2} \in (0, 1)$ as in [Fig. 5c](#), where $\kappa_j \rightarrow 0$ induces $\beta_j \rightarrow 0$. Using a user-adjustable cutoff value, many coefficients can be shrunk to zero, enabling the selection of only a few associated omics features (with nonzero coefficients).

Although Bayesian models can quantify uncertainty of estimators more straightforward than penalized regressions, most Bayesian Cox-type models for high-dimensional covariates do not have user-friendly and standalone R packages on CRAN or GitHub. The main reason is the high computational cost of running a high-dimensional Bayesian Cox model. Advanced users with programming capabilities can contact the corresponding authors for original scripts. Since Bayesian priors are more flexible than frequentist Lasso-type penalties, it can be easier to extend Bayesian models by changing shrinkage priors while keeping almost the same algorithm framework. This means that the Bayesian framework can be more suitable if one is interested in tailoring the shrinkage effects, e.g. to include prior knowledge about the importance of omics features, e.g. features corresponding to a molecular pathway that is known to be affected in the disease under study. For more information on different Bayesian priors in cancer prognosis, we suggest a recent review by [Chu *et al.* \(2022\)](#) which summarized other different shrinkage priors on regression coefficients, such as Gaussian-




Figure 5. Density of shrinkage priors in Bayesian survival modeling. (a) Density of the double exponential (Laplace) prior. (b) Density of the mixture spike-and-slab prior. The spike component $\delta_0(\beta_j)$ induces $\beta_j = 0$ and the slab component $\mathcal{N}(0, \tau_j^2)$ induces $\beta_j \neq 0$. (c) Density of the horseshoe prior. The shrinkage weight κ_j close to 0 shrinks β_j toward zero, and κ_j close to 1 allows β_j to escape the shrinkage.

gamma, Gaussian, Cauchy, pMOM (product moment distribution), piMOM (product inverse moment distribution) and peNMIG (parameter-expanded normal-mixture-of-inverse-gamma distribution) priors.

5 Survival model validation

Model validation plays an important role in identifying potential issues, such as model misspecification or overfitting. This is achieved by revisiting the model's specifications and assumptions following model estimation. For example, the Cox model (1) requires proportional hazards and the logarithm of the hazard to be linear with respect to the model covariates. The former assumption can for example be checked by the cumulative Schoenfeld residuals (Grambsch and Therneau 1994), and the latter assumption can be checked by plotting a nonlinear functional form (e.g. spline) for the effect of a covariate. If the Cox model assumptions are not satisfied, one can try certain transformations of covariates (e.g. Box-Cox power transformations, Box and Cox 1964), allow time-varying coefficients or model interactions among covariates (Ng *et al.* 2023), or investigate patient stratification using unsupervised approaches (Cristescu *et al.* 2015). However, the assumption checks are usually suitable only for low-dimensional models, i.e. for a few clinical variables or a few factors projected from the high-dimensional omics feature space. Novel approaches for assumption checks in general high-dimensional settings require further methodological developments. Johnson and Long (2011) used heuristic methods to investigate the Cox model assumptions by separately fitting univariate Cox models one feature at one time, and check P -values for the score tests of individual features and P -values for testing the proportional hazards assumption of univariate Cox models. But the univariate Cox models do not take into account confounding variables. An alternative approach is to loosen the model assumptions for a more robust modeling approach. One example developed specifically for feature selection under possibly nonproportional hazards in a high-dimensional space is concordance regression (Dunkler *et al.* 2010).

5.1 Feature stability analysis

One important aspect in model validation when using omics or other high-dimensional data is the potential instability of feature selection (Kalousis *et al.* 2007). Feature selection using penalized regressions as described in Section 4.2 heavily depends on the values of the penalty parameters [e.g. for the λ parameter in Lasso Cox model (2)]. The penalty parameters

are often optimized by cross-validation (CV) or other resampling methods, and the uncertainty associated with the random selection of subsets may result in uncertainty in the feature selection, e.g. different CV folds will typically result in different selected features. A straightforward way to identify the most stable features is to find the overlap of identified omics features among different data subsets (e.g. CV folds or resamples) to avoid high false discovery rate (Zucknick *et al.* 2008). One can also perform stability selection (Meinshausen and Bühlmann 2010), which allows to select the most stable features at a given Type-I error level for a Lasso or Elastic Net Cox model (Sill *et al.* 2014).

For the Bayesian models in Section 4.3, feature selection stability is naturally assessed by the uncertainty of coefficients' estimators, as reflected in the posterior variances of β_j or the posterior selection probabilities $p(\gamma_j|\mathcal{D})$ (in SSVS), which is a natural benefit of utilizing full Bayesian inference. Although the uncertainty in feature selection introduces increased variability in the predicted survival probabilities, in the Bayesian framework, this can be addressed quite naturally by averaging the survival predictions over all models using Bayesian model averaging (Volinsky *et al.* 1997). If one is interested in a single model, rather than model averaging, the median probability model (Barbieri and Berger 2004) can be used for uncertainty analyses in survival and high-dimensional omics data (Madjar *et al.* 2021).

5.2 Survival prediction and calibration

The fundamental goal of any statistical prediction model is to achieve a better prediction performance than an existing statistical model which we could call the 'conventional model' (Gerds and Kattan 2021). The special nature of the combination of clinical and/or other known prognostic factors (typically low-dimensional, with established effect) and novel omics features (high-dimensional, with unknown effect) should especially be taken into account. Therefore, a survival model consisting of only established clinical and/or other known prognostic factors should serve as the benchmark (i.e. conventional model) for the upcoming modeling. The inclusion of new covariates (e.g. omics features) into a prognostic model only makes sense if the new covariates have added prognostic value over the established clinical prognostic factors (De Bin *et al.* 2014), i.e. the new prognostic model consisting of the new covariates plus benchmark covariates improves the prediction performance over the conventional model.

To confirm that the identified clinical and omics features have prognostic power with respect to the prediction of

patients' survival outcomes, a model should be both accurate (low prediction error) and precise (low prediction uncertainty). The simplest way to demonstrate the prognostic power is to dichotomize the prognostic scores [i.e. linear predictor $X_i\beta$ in the Cox model, Cox (1957)] by its median value, and then use a log-rank test to compare the survival probabilities of the patients in the two groups, see Fig. 2b. Similarly, one can categorize the prognostic scores by multiple quantiles (e.g. 25%, 50% median, and 75%) into multiple groups of patients and perform a log-rank test.

To validate a prediction model systematically (Rahman *et al.* 2017, Royston and Altman 2013), the predictive performance of the model is commonly addressed by

- **discrimination:** the ability of the model to distinguish between low and high risk patients,

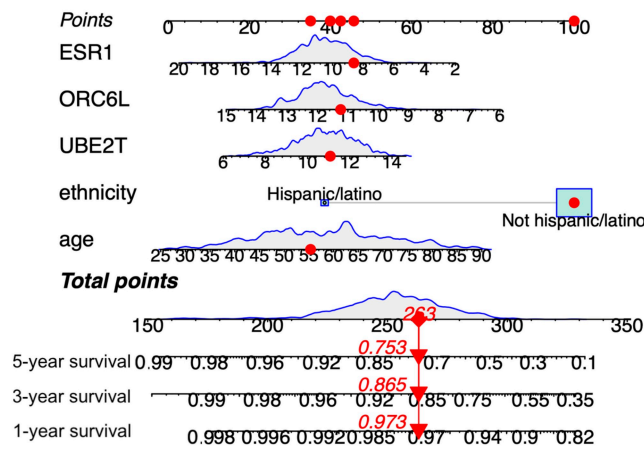


Figure 6. Nomogram developed to estimate the overall survival probability for TCGA-BRCA patients based on clinical (age) demographic (ethnicity) and three selected mRNA features form a Lasso Cox model. The dot symbols represent example patient's information and the triangle symbols represent predicted probabilities of 1-, 3-, and 5-year survival.

- **calibration:** the agreement between the observed outcomes and predicted survival probabilities, and
- **overall performance:** the distance between the observed and predicted survival probabilities.

5.2.1 Discrimination performance

If one focuses on survival prediction at a fixed time point (e.g. 5-year survival probability), a receiver operating characteristic (ROC) curve can be used to evaluate the prognostic (i.e. prediction or discrimination of survival) ability of the survival model, often summarized by its area under the ROC curve (AUC) (Heagerty *et al.* 2000), see Fig. 7a for an example. An AUC of 0.5 is equivalent to the toss of a coin, and the closer the AUC is to 1, the more predictive is the model. When making predictions at multiple time points, ROC curves can be summarized as time-dependent AUC scores, i.e. AUC scores calculated at prespecified time points. Alternatively, the concordance index (C-index) provides a more global, time-independent assessment of the discrimination ability of a prognostic model, such that a better model predicts higher prognostic scores for patients with shorter survival times (Harrell *et al.* 1982, Antolini *et al.* 2005), i.e.

$$C = \mathbb{P}\{S(t|X_i(t)) < S(t|X_j(t)) | T_i < T_j \ \& \ \delta_i = 1\},$$

which means in the absence of censoring, any pair of individuals $\{i, j\}$ with survival times $T_i < T_j$ is concordant if and only if $S(t|X_i(t)) < S(t|X_j(t))$ for any t (equivalent to ranking the prognostic scores $X_i\beta > X_j\beta$ in a Cox model), where t denotes the time instants where there are covariate variations. The C-index can be expressed as a weighted average of the time-dependent AUC over time (Heagerty and Zheng 2005). Therefore, its interpretation is similar to the AUC, where a C-index of 0.5 indicates random predictions, while a perfect prognostic model would have a C-index of 1. There are multiple types of C-indices for survival modeling, in particular the most frequently used Harrell's (Harrell *et al.* 1982) and Uno's C-index (Uno *et al.* 2011). Uno's C-index is more

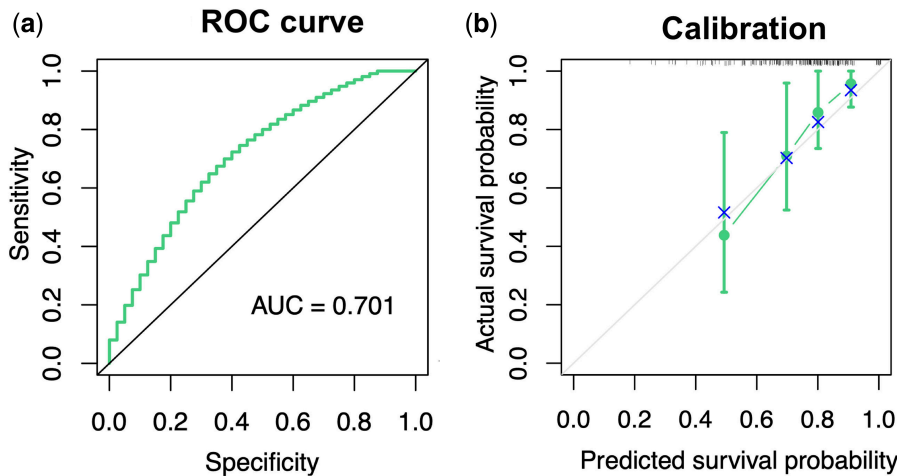


Figure 7. Discrimination and calibration of survival prediction for 20% TCGA-BRCA validation data. TCGA-BRCA data were split to a 80% training set and a 20% validation set. (a) Receiver operating characteristic (ROC) curve estimated at 5-year survival evaluation time point based on the Cox model in Fig. 6. The AUC score is the area under the ROC curve. The 45-degree line represents the performance of a random prediction of the outcome event with AUC = 0.5. (b) Calibration plot estimating 5-year survival probabilities in TCGA-BRCA patients grouped by the quartiles of the predicted probabilities based on the Cox model in Fig. 6. The actual 5-year survival probabilities in each group were estimated by the Kaplan–Meier (KM) method with a 95% confidence interval by bootstrapping. The cross symbol indicates a bias-corrected KM estimate. The predicted survival probability of each quartile group was the mean of the 5-year survival probabilities based on the Cox model for the corresponding group.

robust than Harrell's C-index, in case there is dependence of the concordance on the study-specific censoring distribution.

In the classical Cox modeling framework, both Harrell's and Uno's C-indices only depend on the linear predictors $\mathbf{X}_i\boldsymbol{\beta}$, which is independent of t . But if a model includes covariates with time-dependent effects $\boldsymbol{\beta}(t)$ and/or time-dependent covariates $\mathbf{X}_i(t)$, Harrell's and Uno's C-indices are difficult to be calculated, since they require the calculation of survival functions for each individual over time. In this context, [Antolini et al. \(2005\)](#) proposed a time-dependent C-index, which assesses the concordance of a model's survival distribution predictions $\hat{S}(t)$. This means that Antolini's C-index requires the full specification of $\hat{S}(t)$, even though a C-index only compares the survival probabilities between any pair of individuals, that is, it only assesses whether the relative order of estimated survival probabilities is concordant with observed survival outcomes ([Blanche et al. 2019](#)). Time-dependent prediction indices can better evaluate a model including candidate features with time-dependent effects and/or time-dependent features. To avoid C-hacking among different C-indices in model comparison, [Sonabend et al. \(2022\)](#) recommended that if all models make survival distribution predictions, then select a time-dependent C-index; otherwise choose a time-independent measure (e.g. Uno's C-index); if there is a combination of risk- and distribution-predicting models, then choose a transformation method for analysis (e.g. expected mortality).

5.2.2 Calibration performance

Calibration is to quantify the agreement between the observed and predicted outcomes, which is useful for both internal and external model validation and is recommended to report routinely. The calibration slope is commonly used ([van Houwelingen 2000](#)), which is the slope of the regression of the observed/actual survival probabilities on the model-predicted survival probabilities. A survival model can be reported with the estimated t -year survival probabilities in predefined subgroups, denoted as $S_{\text{model}}(t|g)$ for subgroups $g = 1, \dots, G$. The observed t -year survival probabilities in the subgroups can be estimated by the KM method, denoted as $S_{\text{KM}}(t|g)$. Using the $\ln(-\ln(\cdot))$ -link, the calibration model is

$$\ln(-\ln(S_{\text{KM}}(t|g))) = \alpha + \beta \ln(-\ln(S_{\text{model}}(t|g))) + \epsilon, \quad (6)$$

where ϵ is an error term. If the intercept $\alpha = 0$ and the slope $\beta = 1$, it means that the survival prediction model is well calibrated. For example, [Fig. 7b](#) shows a calibration plot, visualizing the calibration of the estimated 5-year survival probabilities (with 95% confidence interval by bootstrapping) using the KM method for TCGA-BRCA patients grouped by the quartiles of Cox-model predicted survival probabilities. Furthermore, one can calibrate a Cox model in terms of the baseline cumulative hazard and prognostic score. For nonproportional hazard models, calibration using the model cumulative hazard function can be considered ([van Houwelingen 2000](#)).

As an alternative to the calibration slope at a single time point, [Andres et al. \(2018\)](#) and [Haider et al. \(2020\)](#) suggested the distributional calibration (D-Calibration) for accounting survival probabilities across all time points. This can be useful when assessing the entire post-treatment survival prediction, e.g. assessing post liver transplantation survival utility in [Andres et al. \(2018\)](#).

5.2.3 Overall performance

Scoring rules can evaluate the accuracy and confidence of probabilistic predictions, and assess both discrimination and calibration ([Gneiting and Raftery 2007](#), [Avati et al. 2020](#)). The idea of scoring rules dates back to [Brier \(1950\)](#) which assigned a numerical score for verifying ensemble-based probabilistic predictions of discrete outcomes.

[Graf et al. \(1999\)](#) proposed the time-dependent Brier score, which is the expected mean-squared error of survival probability prediction at different time points, i.e.

$$\text{BS}(t) = \frac{1}{n} \sum_{i=1}^n \left\{ \frac{\hat{S}(t|\mathbf{X}_i)^2 \mathbb{1}\{t_i \leq t \ \& \ \delta_i = 1\}}{\hat{G}(t_i)} + \frac{\{1 - \hat{S}(t|\mathbf{X}_i)\}^2 \mathbb{1}\{t_i > t\}}{\hat{G}(t)} \right\},$$

where t_i is the survival time of i th individual, $\hat{S}(t|\mathbf{X}_i)$ is the Cox-model predicted survival probability and $\hat{G}(t)$ is the KM estimate of the censoring distribution. The benefit of the Brier score is that it does not only measure discrimination, similar to evaluation measures like the C-index, but also calibration performance of a survival model. The integrated Brier score (IBS) is as a single measure of prediction accuracy integrating $\text{BS}(t)$ over an entire follow-up time period. [Hielscher et al. \(2010\)](#) presented a comparison between the IBS and a D measure ([Schemper and Henderson 2000](#)), which is an integrated measure based on the mean absolute deviation rather than the mean-squared error used in IBS. The D measure is more robust toward extreme observations and has a smaller variance than the IBS.

To overcome potential overfitting when using feature selection and model estimation, the survival predictions and model calibration should be evaluated in independent validation datasets. As independent validation data are seldom available, we can split the available data into training and validation datasets. However, any single split of the data hides partial data from all steps of model building, which might introduce bias, it is recommended to use resampling-based methods for assessing the survival model's performance. In addition, using resampling-based methods also allow us to estimate the uncertainty of the performance estimator ([Sill et al. 2014](#); [Gerds and Kattan 2021](#), chapter 7). This can be done for example by repeatedly splitting the data to training and validation sets, and evaluating a survival model's performance on the different validation sets using various discrimination or calibration indices. The 0.632+ bootstrap estimator for a discrimination or calibration index can balance the apparent (training) error and the average out-of-bag bootstrap error, and in addition accounts for the relative overfitting based on a non-information error rate in high-dimensional settings ([Schumacher et al. 2007](#), [Binder and Schumacher 2008a](#)). This is a typical machine learning approach with two levels of resampling. The outer layer of resampling is to evaluate the prediction performance and the inner layer of resampling (usually CV) is to optimize model's tuning parameters. We note that the preparatory steps, such as multimodal data standardization and feature pre-selection in the context of high-dimensional input data, may affect the survival prediction performance, and thus should be included in the resampling steps for model validation.

5.2.4 Graphical representation

After confirming that a model is valid (assumptions hold), accurate (low prediction error), precise (uncertainty of performance measures properly quantified) and its predictions generalizable beyond the training dataset (using independent validation data if available), a prognostic nomogram (Kattan *et al.* 1999) can be used to summarize the prognostic effect of the identified clinical and omics features on the risk of a specific year's overall survival (Fig. 6), which may help the clinicians to enhance the patient management and personalized treatment strategies. For example, the red colored dots in Fig. 6 show the information of the identified five variables from an example patient and the corresponding scoring points. The summed scoring points of 263 maps to the predicted 1-year, 3-year and 5-year survival probabilities of this patient. Note that most nomograms treat the identified variables independently in the risk calculation, even though there may be significant interactions among the model features that were used in the feature selection step. However, visualizing such interaction effects would make the nomograms less accessible and interpretable, and so, there is still a room for improvement in how to translate the multivariate risk scores into clinical practice.

When an independent validation dataset is available, it is recommended to report a calibration plot corresponding to the nomogram. Using independent validation data to obtain $S_{\text{model}}(t|g)$ in the calibration model (6) is for the generalization capacity of the model. Since we here do not have independent validation data besides TCGA-BRCA data, Fig. 7 shows an example calibration plot at 5-year survival evaluation time point based on the built Cox model in Fig. 6 for a split 20% TCGA-BRCA dataset.

6 Beyond penalized and Bayesian Cox models

In this tutorial, we mainly focused on penalized regressions and Bayesian hierarchical models in the Cox proportional hazards framework. One can extend this framework in several ways. For instance, one can stay in a likelihood-based modeling framework, but replace the partial likelihood function of the semi-parametric Cox model by alternative likelihood functions (which do not necessarily need to imply proportional hazards), e.g. parametric survival models like exponential, Weibull, or accelerated failure time (AFT) models, or Aalen's additive hazard model (Gorst-Rasmussen and Scheike 2012). Alternatively, one can move to a more algorithmic machine learning approach, such as tree-based boosting or bagging methods, e.g. random survival forests (Hothorn *et al.* 2006, Binder and Schumacher 2008b, Jaeger *et al.* 2019), or (deep) neural networks (Wiegrebe *et al.* 2024).

Hothorn *et al.* (2006) introduced ensemble tree methods for analyzing right-censored survival data, which construct ensembles from base learners, e.g. binary survival trees for each omics feature. Hothorn *et al.* (2006) also proposed a gradient boosting algorithm to predict the survival time of patients with acute myeloid leukemia (AML), based on clinical and omics features. Similarly, Binder and Schumacher (2008b) developed a likelihood-based boosting method, which aims to maximize the Cox partial likelihood function, for modeling time-to-event data based on high-dimensional omics input data and which also allows the inclusion of a small number of mandatory covariates. In general, one needs

to be cautious if using some machine learning methods that are not well-suited for high-dimensional features. For example, Kvamme *et al.* (2019) proposed extensions of the Cox model with neural networks, which are only valid if the number of covariates is smaller than the number of samples, i.e. if $p < n$. A systematic review of deep learning for survival analysis, which includes a survey of methods suitable for high-dimensional data ($p > n$), is provided by Wiegrebe *et al.* (2024).

In the case of nonproportional hazards, many likelihood-based survival models beyond Cox-type models have also been extended to account for high-dimensional omics as input data. For example, Ma *et al.* (2006) combined Lin and Yin's additive hazard model (Lin and Ying 1994) with principal component regression for dimension reduction of omics features, which was applied to the study of gene expression-based survival prediction for diffuse large B-cell lymphoma. Engler and Li (2009) added the Elastic Net penalty in an accelerated ATF model, which assumes that the effect of a covariate accelerates or decelerates the life course of patients. Schmid and Hothorn (2008) and Barnwal *et al.* (2022) used boosting algorithms to learn parametric AFT models. Ha *et al.* (2014) considered the Lasso, SCAD and a penalized likelihood for feature selection in frailty models which assume that individuals have unobserved heterogeneity captured by a latent random term Z , which adapts the Cox model (1) into $h(t|X) = Zh_0(t) \exp(X\beta)$.

6.1 Advanced survival models: cure models, competing risks and multi-state models

In some situations, survival data may be different from Fig. 1c (also Section 2.1), where it was presumed that all individuals will eventually experience the event of interest. Liu *et al.* (2012) studied the Lasso and SCAD feature selection for the proportional hazard mixture cure model, in which a certain fraction of individuals will never experience the event of interest. Tapak *et al.* (2015) investigated Lasso, Elastic Net and likelihood-based boosting for microarray-based survival modeling with competing risks, such as 'progression' versus 'death from noncancer cause', i.e. the event of a patient can occur due to one of multiple distinct causes. There is a growing awareness of the impact of competing risks when developing prognostic models with high-dimensional input data, e.g. Binder *et al.* (2009), Ambrogi and Scheike (2016) and Fu *et al.* (2017). For a single individual who can experience several possible events, Dutta *et al.* (2017) proposed a multi-state model to identify risk factors in different stages of disease based on high-dimensional input data.

7 Toward single-cell data analysis

The cellular heterogeneity of complex sample mixtures pose challenges and also opportunities for precision medicine and survival prediction. For example, Zhou *et al.* (2019) showed that tumor microenvironment-related gene expression signatures do not only accurately predict the survival among colon cancer patients, but also serve as biomarkers for identifying patients who could potentially benefit from adjuvant chemotherapy. Single-cell technologies provide an unprecedented opportunity for dissecting the interplay between the cancer cells and the associated tumor microenvironment, and the produced high-dimensional data should also augment existing survival modeling approaches. The emerging single-cell

atlases are providing a detailed and quantitative overview of tissue composition and organization, and will advance both biomedical research and clinical practice (Elmentaite *et al.* 2022).

Currently, the focus of statistical model development for single-cell data analysis is to understand cell type composition and its impact on gene regulation and transcriptional dynamics, usually based on only a small number of samples/individuals. On the one hand, the underlying statistical models for single-cell data analysis are still in development and continuously being re-evaluated and challenged (Kharchenko 2021). On the other hand, survival analysis tackles disease and health contexts at the individual level, and usually requires a relatively large number of individuals for sufficient statistical power. In particular, the development of improved computational methods is urgently needed to enable the consideration of multiple layers (e.g. individual-level, cellular-level, molecular-level) of information when integrating groups of individuals and omics data from a variety of molecular modalities. Therefore, current approaches often map the findings from single-cell omics data to large-scale bulk sequencing omics and survival data, rather than jointly analyzing single-cell omics and survival data. For example, Guo *et al.* (2018) introduced a generalizable approach to first study T-cells in 14 nonsmall cell lung cancer (NSCLC) patients and to identify gene signatures from the tumor-enriched T cell clusters, and then investigated these gene signatures using bulk RNA-seq and survival data from larger TCGA-NSCLC patient cohort. Similarly, Zhang *et al.* (2020) developed a scRNA-seq-based approach to reconstruct a multilayer signaling network based on 16 glioma patients, and then investigated the network genes using survival data from TCGA Chinese glioma genome atlas (TCGA-CGGA) patients. However, direct joint analysis of survival and single-cell omics data from multiple cellular hierarchies requires further methodological developments and new statistical and machine learning methods.

8 Discussion

Although survival analysis faces many modeling challenges, mainly due to censored outcomes, it represents a well-established methodology for finding risk factors associated with patients' survival. The identification of omics biomarkers for survival prognosis may provide systematic means to guide patient management and personalized treatment and diagnostic strategies. In this tutorial, we provided a comprehensive workflow for survival analysis with high-dimensional omics and standard clinical data, with a specific focus on feature selection of survival-associated omics features and survival model validation. We covered many penalized regressions and Bayesian models for feature selection and survival prediction, accounting for their specific assumptions and applications. Examples of real data and R scripts have been made available to illustrate the use of different methods, which should help researchers to choose and apply suitable methods for their survival analysis applications (<https://ocbeuio.github.io/survomics>). We note that this review only considers methods for right-censored time-to-event data, i.e. where all individuals are assumed to be followed continuously from the start of the study, but where the follow-up period might end before the event (e.g. death) was observed. Other types of censoring include interval censoring and left

truncation, and appropriate statistical methods dealing with these censoring patterns should be chosen accordingly.

Most of the current methods for survival analysis do not explicitly take into account the complex structures within and between multi-omics data, such as gene regulation and DNA-protein interactions. Regulatory networks constructed either based on prior biological knowledge or using data-driven, yet biologically explainable approaches, may help establish useful methodologies for survival analysis that are more effective for deriving biological insights as well as enable improved clinical translation. However, to achieve a comprehensive and biologically meaningful integration of high-dimensional multi-omics data, there is a need for continued development of computational and statistical approaches that consider both technical and biological intricacies of the data and technologies, respectively (Wissel *et al.* 2023). This is currently a very active research field, and we expect to see many improved multi-omics methods for survival prediction in the future.

Another limitation of most of the reviewed methods is that they identify omics features prognostic of survival, but they cannot determine causal relationships. Causality is a fundamental notion to understand omics features causing disease progression, which will allow one to reliably intervene omics features for targeted therapies. There are two popular causal inference models, Pearl's structural causal model (SCM) and Rubin's causal model (RCM), both of which introduce perturbations to draw causal inference. Farooq *et al.* (2023) utilized SCM-based causal discovery approaches to unravel relationships between omics features and survival of breast cancer patients. However, to identify reliable causal relations for clinical applications, laboratory-based experiments, e.g. clustered regularly interspaced short palindromic repeats (CRISPR) techniques (Wang and Doudna 2023), are often necessary to verify the functional relevance of the identified omics features. High-dimensional RCM-based mediation analysis has been used to investigate the indirect effect transmitted by omics features between an exposure and survival outcomes (Song *et al.* 2020, 2021). Causal mediation analysis is an important tool, which considers the problem of decomposing the causal effect of treatment/exposure into direct and indirect effects (Lange and Hansen 2011, VanderWeele 2011). The direct effect corresponds to the effect of a treatment directly on the survival outcome, while an indirect effect corresponds to the effect of a treatment on the outcome that is due to its effect on an intermediate variable (e.g. gene expression) that also has a causal effect on the survival outcome. Targeted learning also fills a much needed gap between statistical modeling and causal inference (van der Laan and Rose 2011, 2018). Tuglus and van der Laan (2011) used targeted maximum likelihood estimation to provide an interpretable causal measure of variable importance for the discovery of biomarkers and omics features. Another way to formalize personalized medicine is dynamic treatment regimes (Chakraborty and Moodie 2013, Tsiatis *et al.* 2019, Deliu *et al.* 2023) that encompasses causal inference and takes into account for the variability in omics, environment and lifestyle factors for each individual to improve the treatment of a particular patient. However, all the causal machine learning methods require further methodological developments for adaptation to survival modeling with high-dimensional input data.

Acknowledgements

The authors thank Professor Ørnulf Borgan for his valuable suggestions and comments.

Conflict of interest

None declared.

Funding

This work was supported by grants from Helse Sør-Øst [2020026 to T.A.], the Norwegian Cancer Society [216104 to T.A.], the Radium Hospital Foundation, the Academy of Finland [326238, 340141, 344698, 345803 to T.A.], Cancer Society of Finland [to T.A.], the European Union's Horizon 2020 research and innovation programme ['PANCAIM' 101016851 to J.Z. and T.A., 'RESCUER' 847912 to M.Z.], the Innovative Medicines Initiative 2 Joint Undertaking of the European Union's Horizon 2020 research and innovation programme and EFPIA and JDRF INTERNATIONAL ['imSAVAR' 853988 to M.Z.], and ERA PerMed under the ERA-NET Cofund scheme of the European Union's Horizon 2020 research and innovation framework programme ['SYMMETRY' 342752 to M.Z.].

Data availability

Supplementary step-by-step R tutorial is available online at <https://ocbe-uo.github.io/survomics>. TCGA data are publicly available at <https://portal.gdc.cancer.gov>.

References

- Aittokallio T. Dealing with missing values in large-scale studies: microarray data imputation and beyond. *Brief Bioinform* 2009; 11:253–64.
- Akbani R, Ng PKS, Werner HMJ *et al.* A pan-cancer proteomic perspective on the cancer genome atlas. *Nat Commun* 2014;5:3887.
- Ambroggi F, Scheike TH. Penalized estimation for competing risks regression with applications to high-dimensional covariates. *Biostatistics* 2016;17:708–21.
- Andres A, Montano-Loza A, Greiner R *et al.* A novel learning algorithm to predict individual survival after liver transplantation for primary sclerosing cholangitis. *PLoS One* 2018;13:e0193523.
- Antolini L, Boracchi P, Biganzoli E. A time-dependent discrimination index for survival data. *Stat Med* 2005;24:3927–44.
- Avati A, Duan T, Zhou S *et al.* Countdown regression: sharp and calibrated survival predictions. In: Adams RP, Gogate V (eds.), *Proceedings of The 35th Uncertainty in Artificial Intelligence Conference*, volume 115 of *Proceedings of Machine Learning Research* (PMLR). Tel Aviv, Israel: ML Research Press. 2020, 145–55.
- Bair E, Tibshirani R. Semi-supervised methods to predict patient survival from gene expression data. *PLoS Biol* 2004;2:e108.
- Barbieri MM, Berger JO. Optimal predictive model selection. *Ann. Stat* 2004;32:870–97.
- Barnwal A, Cho H, Hocking T. Survival regression with accelerated failure time model in XGBoost. *J. Comput. Graph Stat* 2022; 31:1292–302.
- Bartel DP. Metazoan MicroRNAs. *Cell* 2018;173:20–51.
- Binder H, Allignol A, Schumacher M *et al.* Boosting for high-dimensional time-to-event data with competing risks. *Bioinformatics* 2009;25:890–6.
- Binder H, Schumacher M. Adapting prediction error estimates for biased complexity selection in high-dimensional bootstrap samples. *Stat Appl Genet Mol* 2008a;7:1.
- Binder H, Schumacher M. Allowing for mandatory covariates in boosting estimation of sparse high-dimensional survival models. *BMC Bioinformatics* 2008b;9:14.
- Blanche P, Kattan MW, Gerds TA. The c-index is not proper for the evaluation of t-year predicted risks. *Biostatistics* 2019;20:347–57.
- Bommert A, Welchowski T, Schmid M *et al.* Benchmark of filter methods for feature selection in high-dimensional gene expression survival data. *Brief. Bioinform* 2022;23:bbab354.
- Bøvelstad HM, Nygård S, Størvold HL *et al.* Predicting survival from microarray data – a comparative study. *Bioinformatics* 2007; 23:2080–7.
- Bøvelstad HM, Nygård S, Borgan O. Survival prediction from clinico-genomic models – a comparative study. *BMC Bioinformatics* 2009; 10:413.
- Box GEP, Cox DR. An analysis of transformations. *J R Stat Soc B Met* 1964;26:211–52.
- Bradburn MJ, Clark TG, Love SB *et al.* Survival analysis part II: multivariate data analysis – an introduction to concepts and methods. *Br J Cancer* 2003;89:431–6.
- Brier GW. Verification of forecasts expressed in terms of probability. *Mon Weather Rev* 1950;78:1–3.
- Cairns RA, Harris IS, Mak TW. Regulation of cancer cell metabolism. *Nat Rev Cancer* 2011;11:85–95.
- Carvalho CM, Polson NG, Scott JG. Handling sparsity via the horseshoe. In: van Dyk D, Welling M (eds.), *Proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics*, volume 5 of *Proceedings of Machine Learning Research*, Hilton Clearwater Beach Resort, Clearwater Beach, Florida USA. PMLR. 2009, 73–80.
- Chakraborty B, Moodie EE. *Statistical Methods for Dynamic Treatment Regimes: Reinforcement Learning, Causal Inference and Personalized Medicine*. New York: Springer, 2013.
- Chekouo T, Stingo FC, Doecke JD *et al.* miRNA-target gene regulatory networks: a Bayesian integrative approach to biomarker selection with application to kidney cancer. *Biometrics* 2015;71:428–38.
- Chu J, Sun NA, Hu W *et al.* The application of Bayesian methods in cancer prognosis and prediction. *Cancer Genomics Proteomics* 2022;19:1–11.
- Clark TG, Bradburn MJ, Love SB *et al.* Survival analysis part IV: further concepts and methods in survival analysis. *Br J Cancer* 2003; 89:781–6.
- Cox DR. Note on grouping. *J Am Stat Assoc* 1957;52:543–7.
- Cox DR. Regression models and life-tables. *J R Stat Soc B Met* 1972; 34:187–202.
- Cristescu R, Lee J, Nebozhyn M *et al.* Molecular analysis of gastric cancer identifies subtypes associated with distinct clinical outcomes. *Nat Med* 2015;21:449–56.
- De Bin R, Sauerbrei W, Boulesteix A-L. Investigating the prediction ability of survival models based on both clinical and omics data: two case studies. *Stat Med* 2014;33:5310–29.
- de Brevern A, Hazout S, Malpertuy A. Influence of microarrays experiments missing values on the stability of gene groups by hierarchical clustering. *BMC Bioinformatics* 2004;5:114.
- Deliu N, Williams JJ, Chakraborty B. Reinforcement learning in modern biostatistics: Constructing optimal adaptive interventions. arXiv, arXiv:2203.02605, 2023, preprint: not peer reviewed.
- Dunkler D, Schemper M, Heinze G. Gene selection in microarray survival studies under possibly non-proportional hazards. *Bioinformatics* 2010;26:784–90.
- Dutta S, Datta S, Datta S. Temporal prediction of future state occupation in a multistate model from high-dimensional baseline covariates via pseudo-value regression. *J Stat Comput Simul* 2017; 87:1363–78.
- Edwards NJ, Oberti M, Thangudu RR *et al.* The CPTAC data portal: a resource for cancer proteomics research. *J Proteome Res* 2015; 14:2707–13.
- Elmentaite R, Domínguez Conde C, Yang L *et al.* Single-cell atlases: shared and tissue-specific cell types across human organs. *Nat Rev Genet* 2022;23:395–410.

- Engler D, Li Y. Survival analysis with high-dimensional covariates: an application in microarray studies. *Stat Appl Genet Mol* 2009; 8:1–22.
- Fan J, Feng Y, Wu Y. High-dimensional variable selection for Cox's proportional hazards model. *IMS Collections* 2010;6:70–86.
- Fan J, Li R. Variable selection for Cox's proportional hazards model and frailty model. *Ann. Stat* 2002;30:74–99.
- Farooq M, Hardan S, Zhumbhayeva A *et al.* Understanding breast cancer survival: using causality and language models on multi-omics data. arXiv, arXiv:2305.18410, 2023, preprint: not peer reviewed.
- Fu Z, Parikh CR, Zhou B. Penalized variable selection in competing risks regression. *Lifetime Data Anal* 2017;23:353–76.
- George E, McCulloch R. Variable selection via Gibbs sampling. *J Am Stat Assoc* 1993;88:881–9.
- Gerds TA, Kattan MW. *Medical Risk Prediction Models: With Ties to Machine Learning*. Boca Raton, FL: Chapman and Hall/CRC, 2021.
- Gneiting T, Raftery AE. Strictly proper scoring rules, prediction, and estimation. *J Am Stat Assoc* 2007;102:359–78.
- Gorst-Rasmussen A, Scheike TH. Coordinate descent methods for the penalized semiparametric additive hazards model. *J Stat Soft* 2012; 47:1–17.
- Graf E, Schmoor C, Sauerbrei W *et al.* Assessment and comparison of prognostic classification schemes for survival data. *Statist Med* 1999;18:2529–45.
- Grambsch PM, Therneau TM. Proportional hazards tests and diagnostics based on weighted residuals. *Biometrika* 1994;81:515–26.
- Green MS, Symons MJ. A comparison of the logistic risk function and the proportional hazards model in prospective epidemiologic studies. *J Chronic Dis* 1983;36:715–23.
- Guo X, Zhang Y, Zheng L *et al.* Global characterization of T cells in non-small-cell lung cancer by single-cell sequencing. *Nat Med* 2018; 24:978–85.
- Ha ID, Pan J, Oh S *et al.* Variable selection in general frailty models using penalized h-likelihood. *J Comput Graph Stat* 2014;23:1044–60.
- Haider H, Hoehn B, Davis S *et al.* Effective ways to build and evaluate individual survival distributions. *J Mach Learn Res* 2020;21:1–63.
- Harrell FE, Califf RM, Pryor DB *et al.* Evaluating the yield of medical tests. *JAMA* 1982;247:2543–6.
- Hasin Y, Seldin M, Lusis A. Multi-omics approaches to disease. *Genome Biol* 2017;18:83.
- Haukaas TH, Euceda LR, Giskeødegård GF *et al.* Metabolic portraits of breast cancer by HR MAS MR spectroscopy of intact tissue samples. *Metabolites* 2017;7:18.
- Heagerty PJ, Lumley T, Pepe MS. Time-dependent ROC curves for censored survival data and a diagnostic marker. *Biometrics* 2000; 56:337–44.
- Heagerty PJ, Zheng Y. Survival model predictive accuracy and ROC curves. *Biometrics* 2005;61:92–105.
- Heinze G, Wallisch C, Dunkler D. Variable selection – a review and recommendations for the practicing statistician. *Biom J* 2018; 60:431–49.
- Held L, Gravestock I, Sabané Bové D. Objective Bayesian model selection for Cox regression. *Stat Med* 2016;35:5376–90.
- Henderson R. Problems and prediction in survival-data analysis. *Stat Med* 1995;14:161–84.
- Herrmann M, Probst P, Hornung R *et al.* Large-scale benchmark study of survival prediction methods using multi-omics data. *Brief. Bioinform* 2021;22:bbaa167.
- Hielscher T, Zucknick M, Werft W *et al.* On the prognostic value of survival models with application to gene expression signatures. *Stat Med* 2010;29:818–29.
- Hothorn T, Bühlmann P, Dudoit S *et al.* Survival ensembles. *Biostatistics* 2006;7:355–73.
- Hrydziusko O, Viant MR. Missing values in mass spectrometry based metabolomics: an undervalued step in the data processing pipeline. *Metabolomics* 2012;8:161–74.
- Ibrahim JG, Chen M-H, Sinha D. *Bayesian Survival Analysis*, Vol. 2. New York: Springer, 2001.
- Jaeger BC, Long DL, Long DM *et al.* Oblique random survival forests. *Ann Appl Stat* 2019;13:1847–83.
- Jendoubi T. Approaches to integrating metabolomics and multi-omics data: a primer. *Metabolites* 2021;11:184.
- Jiang R, Sun T, Song D *et al.* Statistics or biology: the zero-inflation controversy about scRNA-seq data. *Genome Biol* 2022;23:31.
- Johnson BA, Long Q. Survival ensembles by the sum of pairwise differences with application to lung cancer microarray studies. *Ann Appl Stat* 2011;5:1081–101.
- Kalousis A, Prados J, Hilario M. Stability of feature selection algorithms: a study on high-dimensional spaces. *Knowl Inf Syst* 2007; 12:95–116.
- Kaplan EL, Meier P. Nonparametric estimation from incomplete observations. *J Am Stat Assoc* 1958;53:457–81.
- Kattan MW, Wheeler TM, Scardino PT. Postoperative nomogram for disease recurrence after radical prostatectomy for prostate cancer. *J Clin Oncol* 1999;17:1499–507.
- Kharchenko PV. The triumphs and limitations of computational methods for scRNA-seq. *Nat Methods* 2021;18:723–32.
- Kim J, Sohn I, Jung S-H *et al.* Analysis of survival data with group lasso. *Commun Stat Simul C* 2012;41:1593–605.
- Kong W, Hui HWH, Peng H *et al.* Dealing with missing values in proteomics data. *Proteomics* 2022;22:e2200092.
- Konrath S, Fahrmeir L, Kneib T. Bayesian smoothing, shrinkage and variable selection in hazard regression. In: Becker C, Fried R, Kuhnt S (eds.), *Robustness and Complex Data Structures: Festschrift in Honour of Ursula Gather*, Chapter 10. Berlin: Springer, 2013, 149–70.
- Korthauer K, Kimes PK, Duvallet C *et al.* A practical guide to methods controlling false discoveries in computational biology. *Genome Biol* 2019;20:118.
- Kundu S, Cheng Y, Shin M *et al.* Bayesian variable selection with graphical structure learning: applications in integrative genomics. *PLoS ONE* 2018;13:e0195070.
- Kvamme H, Borgan Ø, Scheel I. Time-to-event prediction with neural networks and Cox regression. *J Mach Learn Res* 2019;20:1–30.
- Lange T, Hansen JV. Direct and indirect effects in a survival context. *Epidemiology* 2011;22:575–81.
- Lee KH, Chakraborty S, Sun J. Bayesian variable selection in semiparametric proportional hazards model for high dimensional survival data. *Int J Biostat* 2011;7:1–32.
- Lee KH, Chakraborty S, Sun J. Survival prediction and variable selection with simultaneous shrinkage and grouping priors. *Stat Anal* 2015;8:114–27.
- Leung K, Elashoff R, Afifi A. Censoring issues in survival analysis. *Annu Rev Public Health* 1997;18:83–104.
- Lin DY, Ying Z. Semiparametric analysis of the additive risk model. *Biometrika* 1994;81:61–71.
- Liu X, Peng Y, Tu D *et al.* Variable selection in semiparametric cure models based on penalized likelihood, with application to breast cancer clinical trials. *Stat Med* 2012;31:2882–91.
- Ma S, Kosorok MR, Fine JP. Additive risk models for survival data with high-dimensional covariates. *Biometrics* 2006;62:202–10.
- Madjar K, Zucknick M, Ickstadt K *et al.* Combining heterogeneous subgroups with graph-structured variable selection priors for Cox regression. *BMC Bioinformatics* 2021;22:586.
- McInnes L, Healy J, Melville J. UMAP: Uniform manifold approximation and projection for dimension reduction. arXiv, arXiv:1802.03426, 2018, preprint: not peer reviewed.
- Meinshausen N, Bühlmann P. Stability selection. *J R Stat Soc B Met* 2010;72:417–73.
- Mu J, Liu Q, Kuo L *et al.* Bayesian variable selection for the Cox regression model with spatially varying coefficients with applications to Louisiana respiratory cancer data. *Biom J* 2021;63:1607–22.
- Nam SY, Jeong J, Lee WK *et al.* Sex-specific effect of body mass index and fasting glucose on gastric cancer risk and all causes mortality: a cohort study of 5.17 million. *Int J Obes (Lond)* 2022;46:1644–51.
- Ng HM, Jiang B, Wong KY. Penalized estimation of a class of single-index varying-coefficient models for integrative genomic analysis. *Biometrical J* 2023;65:2100139.

- Okser S, Pahikkala T, Aittokallio T. Genetic variants and their interactions in disease risk prediction – machine learning and network perspectives. *BioData Min* 2013;**6**:5.
- Perou CM, Sørlie T, Eisen MB *et al.* Molecular portraits of human breast tumours. *Nature* 2000;**406**:747–52.
- Peto R, Peto J. Asymptotically efficient rank invariant test procedures. *J R Stat Soc A Stat* 1972;**135**:185.
- Rahman MS *et al.* Review and evaluation of performance measures for survival prediction models in external validation settings. *BMC Med Res Methodol* 2017;**17**:60.
- Rahnenführer J, De Bin R, Benner A *et al.*; For Topic Group “High-dimensional data” (TG9) of the STRATOS Initiative. Statistical analysis of high-dimensional biomedical data: a gentle introduction to analytical goals, common approaches and challenges. *BMC Med* 2023;**21**:182.
- Royston P, Altman DG. External validation of a cox prognostic model: principles and methods. *BMC Med Res Methodol* 2013;**13**:33.
- Schemper M, Henderson R. Predictive accuracy and explained variation in Cox regression. *Biometrics* 2000;**56**:249–55.
- Schmid M, Hothorn T. Flexible boosting of accelerated failure time models. *BMC Bioinformatics* 2008;**9**:269.
- Schumacher M, Binder H, Gerds T. Assessment of survival prediction models based on microarray data. *Bioinformatics* 2007;**23**:1768–74.
- Sill M, Hielscher T, Becker N *et al.* c060: extended inference with lasso and elastic-net regularized Cox and generalized linear models. *J Stat Soft* 2014;**62**:1–22.
- Simon N, Friedman J, Hastie T *et al.* Regularization paths for Cox’s proportional hazards model via coordinate descent. *J Stat Softw* 2011;**39**:1–13.
- Simon N, Friedman J, Hastie T *et al.* A sparse-group lasso. *J Comput Graph Stat* 2013;**22**:231–45.
- Sonabend R, Bender A, Vollmer S. Avoiding C-hacking when evaluating survival distribution predictions with discrimination measures. *Bioinformatics* 2022;**38**:4178–84.
- Song Y, Zhou X, Zhang M *et al.* Bayesian shrinkage estimation of high dimensional causal mediation effects in omics studies. *Biometrics* 2020;**76**:700–10.
- Song Y, Zhou X, Kang J *et al.* Bayesian hierarchical models for high-dimensional mediation analysis with coordinated selection of correlated mediators. *Stat Med* 2021;**40**:6038–56.
- Sun J, Xia Y. Pretreating and normalizing metabolomics data for statistical analysis. *Genes Dis* 2023;**11**:100979.
- Tapak L, Saidijam M, Sadeghifar M *et al.* Competing risks data analysis with high-dimensional covariates: an application in bladder cancer. *Genomics Proteomics Bioinf* 2015;**13**:169–76.
- Tarazona S, Arzalluz-Luque A, Conesa A. Undisclosed, unmet and neglected challenges in multi-omics studies. *Nat Comput Sci* 2021;**1**:395–402.
- Taylor S, Ponzini M, Wilson M *et al.* Comparison of imputation and imputation-free methods for statistical analysis of mass spectrometry data with missing data. *Brief Bioinform* 2022;**23**:bbab353.
- Tibshirani R. The lasso method for variable selection in the Cox model. *Statist Med* 1997;**16**:385–95.
- Tsiatis AA, Davidian M, Holloway ST *et al.* *Dynamic Treatment Regimes: Statistical Methods for Precision Medicine*. Boca Raton, FL: Chapman and Hall/CRC, 2019.
- Tuglus C, van der Laan MJ. Targeted methods for biomarker discovery. In: *Targeted Learning: Causal Inference for Observational and Experimental Data*. New York: Springer, 2011, 367–82.
- Tuikkala J, Elo L, Nevalainen OS *et al.* Improving missing value estimation in microarray data with gene ontology. *Bioinformatics* 2005;**22**:566–72.
- Uno H, Cai T, Pencina MJ *et al.* On the C-statistics for evaluating overall adequacy of risk prediction procedures with censored survival data. *Stat Med* 2011;**30**:1105–17.
- van der Laan MJ, Rose S. *Targeted Learning: Causal Inference for Observational and Experimental Data*. New York: Springer, 2011.
- van der Laan MJ, Rose S. *Targeted Learning in Data Science: Causal Inference for Complex Longitudinal Studies*. New York: Springer, 2018.
- van der Maaten L, Hinton G. Visualizing data using t-SNE. *J Mach Learn Res* 2008;**9**:2579–605.
- van Houwelingen HC. Validation, calibration, revision and combination of prognostic survival models. *Statist Med* 2000;**19**:3401–15.
- VanderWeele TJ. Causal mediation analysis with survival data. *Epidemiology* 2011;**22**:582–5.
- Vinga S. Structured sparsity regularization for analyzing high-dimensional omics data. *Brief Bioinform* 2021;**22**:77–87.
- Volinsky CT, Madigan D, Raftery AE *et al.* Bayesian model averaging in proportional hazard models: assessing the risk of a stroke. *J R Stat Soc C APP* 1997;**46**:433–48.
- Wang JY, Doudna JA. CRISPR technology: a decade of genome editing is only the beginning. *Science* 2023;**379**:eadd8643.
- Wang W, Liu W. Integration of gene interaction information into a reweighted Lasso-Cox model for accurate survival prediction. *Bioinformatics* 2020;**36**:5405–14.
- Wiegrebe S, Kopper P, Sonabend R *et al.* Deep learning for survival analysis: a review. *Artif Intell Rev* 2024;**57**:65.
- Wissel D, Rowson D, Boeva V. Systematic comparison of multi-omics survival models reveals a widespread lack of noise resistance. *Cell Rep Methods* 2023;**3**:100461.
- Zellner A. On assessing prior distributions and Bayesian regression analysis with g-prior distributions. In: Goel RK, Zellner A (eds.), *Bayesian Inference and Decision Techniques: Essays in Honor of Bruno de Finetti. Studies in Bayesian Econometrics and Statistics*, Vol. 6, Chapter 15. North-Holland: Elsevier, 1986, 223–43.
- Zhang HH, Lu W. Adaptive Lasso for Cox’s proportional hazards model. *Biometrika* 2007;**94**:691–703.
- Zhang J, Guan M, Wang Q *et al.* Single-cell transcriptome-based multi-layer network biomarker for predicting prognosis and therapeutic response of gliomas. *Brief Bioinform* 2020;**21**:1080–97.
- Zhao Z, Banterle M, Lewin A *et al.* Multivariate Bayesian structured variable selection for pharmacogenomic studies. *J R Stat Soc Ser C Appl Stat* 2024;**73**:420–43.
- Zhao Z, Zucknick M. Structured penalized regression for drug sensitivity prediction. *J R Stat Soc Ser C Appl Stat* 2020;**69**:525–45.
- Zhou R, Zeng D, Zhang J *et al.* A robust panel based on tumour micro-environment genes for prognostic prediction and tailoring therapies in stage I-III Colon cancer. *EBioMedicine* 2019;**42**:420–30.
- Zucknick M, Richardson S, Stronach EA. Comparing the characteristics of gene expression profiles derived by univariate and multivariate classification methods. *Stat Appl Genet Mol* 2008;**7**:7.
- Zucknick M, Saadati M, Benner A. Nonidentical twins: comparison of frequentist and Bayesian lasso for Cox models. *Biom J* 2015;**57**:959–81.